

Enhancing the sensitivity of the ATLAS experiment to electroweak supersymmetry

Benjamin Haslum Hodgkinson

Darwin College
University of Cambridge



This thesis is submitted for the degree of
Doctor of Philosophy
July 2023

Abstract

The Standard Model of particle physics describes all observed fundamental particles and their interactions, but is incomplete. The ATLAS experiment at the Large Hadron Collider has performed a suite of searches for new fundamental particles predicted by ‘beyond-the-Standard Model’ theories such as supersymmetry (SUSY). As yet, no discoveries have been made, but a range of simplified SUSY scenarios have been excluded. This thesis presents three approaches for enhancing the sensitivity of the ATLAS experiment to the electroweak production of supersymmetric particles.

First, realistic supersymmetric scenarios contain a large number of free parameters and a rich phenomenology which is not captured by the simplified model interpretations provided by typical ATLAS SUSY searches. To address this, several key electroweak SUSY searches are interpreted in the phenomenological Minimal Supersymmetric Standard Model. The complementarity of these searches with additional constraints from beyond ATLAS, such as limits from dark matter experiments, is also discussed. This provides a comprehensive analysis of the sensitivity of ATLAS to electroweak supersymmetry and highlights scenarios that remain uncovered by current searches and constraints.

Second, a statistical combination of three orthogonal searches for the pair-production of charginos is performed. This extends and deepens the existing sensitivity and exclusion power offered by each analysis alone.

Finally, a machine learning method is presented to improve the reconstruction performance of missing transverse momentum (p_T^{miss}) – an essential variable in many searches for supersymmetry. Two neural-network-based approaches are compared and their potential to enhance the sensitivity of ATLAS to pMSSM scenarios is evaluated.

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text. It is not substantially the same as any work that has already been submitted before for any degree or other qualification except as declared in the preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Benjamin Hodkinson

Acknowledgements

I am privileged to have had the support of many wonderful people throughout my journey in physics. The last four years in particular have been a somewhat tumultuous time, not least because of a global pandemic. I would not have reached the finish line without the following people, to whom I am incredibly grateful.

First, thank you to my supervisor, Tina Potter. I didn't know much about ATLAS or 'experimental' high energy physics when you offered me a PhD position in 2019. Accepting that offer is one of the best decisions I have ever made. I've travelled the world, met incredible people, and learned fascinating things about fundamental physics. Thank you for giving me that opportunity, allowing me the freedom to pursue my interests and for all your support and encouragement over the last four years.

I am also grateful to Sarah Williams, who has supported much of the work in this thesis as an unofficial second supervisor. It has been a pleasure working with you, from the early days of my qualification task to sprinting to the pMSSM finish line. Thank you for being so approachable and for your career advice. Your encouragement and positivity have been an antidote to imposter syndrome.

Thank you also to Holly Pacey for your guidance, encouragement and words of wisdom. Your feedback has made me a better physicist, and your encouragement has helped me realise that I might actually be a half-decent one. You have been and will continue to be, a great role model.

I am grateful to have shared my PhD journey with Dominic Jones and David Rousso. From my first day in the department, to hikes in the Swiss mountains, to finishing our theses, you have been brilliant office-mates, collaborators and friends. It's been cathartic moaning about deadlines, bureaucracy and academic life with both of you. Dom, you have an enviable eye for detail and I've benefited greatly from your knowledge of statistics and BDTs. David, thank you for your chaotic brilliance and sense of humour.

Thank you to the rest of the Cambridge ATLAS group – it’s been a pleasure working with you all. I’ve enjoyed the lively discussions in our weekly meetings and benefited from your feedback on my work. Thank you in particular to Will Fawcett whose experience with the pMSSM has been a great help.

As a member of the largest high-energy physics experiment in the world, my gratitude extends far beyond the borders of Cambridge. I am lucky to have collaborated with many fantastic people from all over the world. Thank you to the ATLAS MET & Pileup Tagging sub-group, which has been a great forum for feedback and development of METNet. Thank you in particular to Max Swiatlowski for your support, encouragement and advice. I would also like to thank Holly Pacey, Michael Holzbock, Matthew Leigh and Johnny Raine for their collaboration in the development of METNet.

The ATLAS pMSSM team has provided lots of helpful feedback and suggestions on my work. In particular, thank you to Brian Petersen for your support and for answering countless technical questions. Thank you also to Jonas Würzinger, Yang Liu, Giordon Stark, Lukas Heinrich, Eric Schanet and Alan Barr.

The highlight of my PhD was undoubtedly my Long Term Attachment at CERN, where I spent a year living in Geneva with a brilliant crowd of fellow students. Thanks to everyone for the ski trips, nights out, weekends in the mountains and all the other fun we had. There are too many members of the ‘LTA gang’ to name, but if you’re reading this, you know who you are.

From my time at Durham, I am also grateful to Jakub Scholtz, Jonas Lindert, Frank Krauss and Gordon Love for guiding my first steps into research.

Outside of the world of physics, thank you to my ‘normal’ friends who have provided welcome distractions over the last few years. It would be unfair to name just a few of you, but insufficient to name none, so in no particular order: Thank you to Josh Keene, Ben ‘Benteke’ Jeffery, Fábio Baptista and Dan Ward for being great friends since my physics education began in secondary school. I was lucky to have a great group of friends at Durham who ensured my time was spent doing much more than just physics – thank you to Joe Ravenhall, Dawn Lockhart, Peter Morgan, Louise Lacourarie, Alice Rouget, Pip Matthews, Maddie Rushforth, Daisy Collins, Jenny Dearsley, Rachel Marrinan, Rachel Scott, Bill Cooper, Chris Kelly, Hugo Heyman and James Bacon. From Darwin College, thank you to Matt Cairns and Nina Reinach for all the pints we had and those we missed due to the pandemic. Thank you also to Sue and co. for your support.

Thank you to my family for always supporting and encouraging me. Thank you to my mum, dad and brother – a good chunk of the work in this thesis was completed in lockdown with them. Thank you especially to my mum and dad for always believing in me and encouraging me through setbacks. Thank you to my grandad for sharing his love of mathematics and puzzles with me from an early age.

Most of all, thank you to Lizzie Strahan. Thank you for occasionally indulging my explanations of supersymmetry or dark matter. Thank you for putting up with me moving to Switzerland for a year. Thank you for reminding me how ridiculous it is to be stressed about objects called ‘winos’ and ‘binos’. Thank you for being there for every high and low of my journey over the last five years, whether it be to celebrate or console. Without you by my side, all of this would have been far more stressful and far less enjoyable, if it would even have been possible at all.

Preface

This thesis describes my research on various aspects of the ATLAS experiment. As one of several thousand members of the ATLAS Collaboration, my work is not performed in isolation, and I have benefited from fruitful collaboration within several ATLAS analysis teams. All diagrams, plots and figures presented in this thesis were produced by me unless otherwise stated in the caption with a reference to the original source. The following outlines my personal contribution to the projects included in this thesis.

- **Reinterpretation of ATLAS Run 2 searches for the electroweak production of supersymmetric particles in the pMSSM**

I am the lead analyser in the effort to interpret Run 2 electroweak SUSY searches in the pMSSM. I performed most aspects of the analysis and produced all results and interpretation. This work is included in Chapter 4 of this thesis which is partly based on the Conference Note in Reference [1]. This project sits within the ATLAS pMSSM analysis group, where pMSSM scans targeting third-generation-squark and general-particle production are also under development. The event generation and model-evaluation workflow outlined in Section 4.2 was developed in collaboration with the pMSSM analysis team. I implemented several of the ATLAS analyses into this workflow and developed the code for producing the final result plots shown in Chapter 4. I also developed the RECAST for the ‘2L0J’ analysis.

- **Statistical combination of searches for chargino production**

As part of the ‘Electroweak SUSY Combinations’ analysis team, I performed statistical combinations of three analyses which targeted a simplified model of chargino production. This work was undertaken in collaboration with Dominic Jones. This work is included in Chapter 5, where the final versions of the plots presented are produced by me. Other members of the analysis team were responsible for statistical combinations of other analyses for other simplified models – those results are not included in this thesis. This work contributed to the Conference Note in Reference [2].

- **METNet: A combined p_T^{miss} working point using a neural network with the ATLAS detector**

The work in Chapter 6 began with my ATLAS authorship task, where I performed optimisation and Monte-Carlo-based performance studies of a neural network which regressed truth-level p_T^{miss} . I was the lead author of and produced all results for, an ATLAS Public Note based on this work, published in Reference [3] (credited to the ATLAS Collaboration). I have also presented the work at two international conferences, the proceedings for which can be found in References [4] and [5]. Chapter 6 is based partly on Reference [3] but contains several developments I have worked on since, in particular a classifier-based approach to METNet which solves several issues related to bias and extrapolation. All results and figures presented in this chapter were produced by me. Similar studies on the regression approach to METNet have also been performed by Matthew Leigh in the context of a Masters dissertation [6] and his own authorship task. The conclusions from Matthew's optimisation studies contributed to the final architecture used to train the regression neural network in Chapter 6.

In addition to the above, in the first year of my PhD, I performed a sensitivity study for a search for direct stau production with a fully leptonic final state. I ultimately concluded that such a search would not have sensitivity with the ATLAS Run 2 dataset. This work is not included in this thesis and was documented internally.

I have also made contributions to the operation and performance of ATLAS. During my 12-month Long-Term Attachment (LTA) at CERN I took shifts in the ATLAS control room on the Data Quality desk. I have also contributed to the ATLAS Jet/EtMiss performance group as MET validation contact. This began in December 2019 as part of my ATLAS authorship task (alongside my work on METNet) and continued until May 2023. I performed validation tasks evaluating the impact of new ATLAS software developments on the performance of missing transverse momentum reconstruction and reported my findings in weekly meetings.

Contents

1. Introduction	1
2. Theoretical background	5
2.1. The Standard Model	5
2.2. Problems with the Standard Model	10
2.2.1. Dark matter	11
2.2.2. The hierarchy problem	12
2.3. Supersymmetry	13
2.3.1. The Minimal Supersymmetric Standard Model	14
2.3.2. Electroweak supersymmetry at the LHC	17
2.3.3. Simplified models	18
2.3.4. The phenomenological MSSM	20
3. Experimental background: From protons to p-values	23
3.1. The Large Hadron Collider	24
3.2. The ATLAS detector	27
3.2.1. Inner detector	28
3.2.2. Calorimeters	29
3.2.3. Muon spectrometer	31
3.2.4. Forward detectors	32
3.3. Triggering	33
3.4. Object reconstruction	33
3.4.1. Low-level objects	33
3.4.2. Electrons and photons	34
3.4.3. Muons	35
3.4.4. Hadronic jets	35
3.4.5. Tau leptons	37
3.4.6. Missing transverse momentum	37

3.5. Analysing ATLAS data	40
3.5.1. Simulated data	40
3.5.2. Event selection	41
3.5.3. Statistical analysis	42
3.6. Reinterpretation tools	47
3.6.1. SIMPLEANALYSIS	48
3.6.2. JSON likelihoods	49
3.6.3. RECAST	49
3.7. Machine learning methods	50
3.7.1. Neural networks	50
3.7.2. Regression	52
3.7.3. Classification	53
4. Interpretation of $\sqrt{s} = 13$ TeV ATLAS searches in the pMSSM	55
4.1. Introduction	55
4.2. Method	57
4.2.1. Overview	57
4.2.2. Scan procedure	59
4.2.3. Analysis approximations	66
4.3. A tale of two scans	69
4.3.1. EWKino scan	69
4.3.2. Bino-DM scan	72
4.3.3. Summary of the scans	74
4.4. The Run 2 ATLAS searches for electroweak supersymmetry	76
4.4.1. Direct searches	77
4.4.2. ATLAS bSM Higgs constraints	98
4.4.3. Overlaps and unique sensitivity	99
4.5. External constraints	103
4.6. Overall ATLAS constraints on electroweakinos	109
4.7. Complementarity with dark matter direct-detection experiments	117
4.7.1. Bino-like LSP	118
4.7.2. Wino- and higgsino-like LSP	118
4.8. Summary & outlook	123
4.8.1. Summary of exclusion	123
4.8.2. Viable benchmark scenarios	123
4.8.3. Outlook	127

5. Statistical combination of searches for chargino production	131
5.1. Introduction	131
5.2. Searches for $\tilde{\chi}_1^+ \tilde{\chi}_1^-$ production	132
5.3. Method	134
5.3.1. Performing a statistical combination	134
5.3.2. Analyses combined for each signal point	135
5.3.3. Correlating systematic uncertainties	136
5.4. Validation	139
5.4.1. Validation of analysis inputs	139
5.4.2. Fit stability	139
5.5. Results	142
5.6. Conclusions	147
6. METNet: A combined p_T^{miss} working point using a neural network	149
6.1. Introduction	149
6.2. Training and testing data	151
6.2.1. Simulated event samples	151
6.2.2. Object definitions	152
6.2.3. Training and testing datasets	153
6.2.4. Event selections	154
6.3. Neural network architecture	157
6.3.1. Input features	157
6.3.2. Common architecture	157
6.3.3. METNet _R architecture	159
6.3.4. METNet _C architecture	161
6.3.5. Summary	163
6.4. Results: Comparison with current p_T^{miss} working points	163
6.4.1. Resolution	163
6.4.2. Bias and tails	165
6.4.3. Signal vs. background separation	167
6.4.4. Discussion: METNet _R vs. METNet _C	170
6.5. Classification scores	172
6.6. Outlook	175
7. Summary and outlook	179
A. Extra combination plots	181

Bibliography	189
List of figures	209
List of tables	223

“Do. Or do not. There is no try.”

— Yoda

Chapter 1.

Introduction

The Standard Model of particle physics describes all observed fundamental particles and their interactions, but is incomplete. Several deficiencies, including the lack of a viable dark matter candidate, motivate an extension to the Standard Model known as *supersymmetry* – the focus of this thesis. The ATLAS experiment is the largest particle detector at the Large Hadron Collider (LHC) at CERN. ATLAS performs both precision measurements of Standard Model particles and searches for new particles predicted by theories such as supersymmetry.

The work contributing to this thesis was completed between 2019 and 2023, a period for the majority of which no new experimental data was collected by ATLAS. In the years prior, the second data-taking run of the Large Hadron Collider (‘Run 2’) produced proton–proton collisions at unprecedented energy and rate, providing the tantalising possibility of a discovery of new physics. Yet, at the commencement of this PhD, most key signature searches for supersymmetry had been completed and produced null results. This leaves the hunt for supersymmetry at a pivot point: a suite of previous searches has excluded a range of key scenarios and the path forward is somewhat unclear.

With no new ATLAS data to analyse, this thesis instead asks the question: how can we extract as much information as possible about supersymmetry from existing, and future, searches? In other words, how can the sensitivity of the ATLAS experiment to supersymmetry be *enhanced*? ‘Enhancement’ of sensitivity is explored through three distinct approaches in the context of searches for the electroweak production of supersymmetric particles.

First, the Run 2 searches were optimised for, and interpreted in so-called *simplified* supersymmetric scenarios. However, simplified models do not capture all of the rich

phenomenology present in more complete supersymmetric models. This makes it difficult to assess the global constraints resulting from the null results of ATLAS searches. In Chapter 4, several key SUSY signature searches are interpreted in the phenomenological Minimal Supersymmetric Standard Model (pMSSM) – a space of more realistic, non-simplified, models of supersymmetry. This enhances our understanding of ATLAS sensitivity by providing the most comprehensive analysis to date of the landscape of electroweak supersymmetry that is/is not covered by ATLAS. Additional constraints from beyond ATLAS are also considered, including dark matter experiments, to analyse viable scenarios that should be targeted with the next generation of searches.

Second, each search for supersymmetry performed with the 2015–2018 ATLAS dataset was evaluated independently. But the statistical power, and therefore sensitivity to new physics, of these searches, could be enhanced if evaluated in combination. In Chapter 5, a statistical combination of three ATLAS searches for the pair-production of charginos (a particle predicted by supersymmetry) is presented. The combination significantly improves the exclusion limits on simplified chargino models compared to each of the analyses alone.

Finally, the sensitivity of a search is partly limited by the detection/measurement accuracy of relevant observables in the ATLAS detector. A key observable in most searches for supersymmetry is missing transverse momentum. In Chapter 6, a machine learning-based estimation of missing transverse momentum, named ‘METNet’, is presented. A neural network is used to combine several distinct ATLAS estimates of missing transverse momentum into a single more accurate estimate which could be used to enhance the sensitivity of future searches for supersymmetry. Simulation-based studies of the behaviour and performance of two iterations of this neural network are presented and recommendations for its use in future searches are discussed. In particular, the potential for METNet to improve sensitivity to pMSSM scenarios from Chapter 4 is evaluated.

Relevant theoretical background information is given in Chapter 2. This includes a summary of the Standard Model and the motivation for supersymmetry as a solution to several problems. The motivation for the electroweak production of supersymmetric particles at energy scales accessible at the LHC is also discussed and the pMSSM is introduced in anticipation of Chapter 4.

Chapter 3 outlines the experimental background relevant to the results in this thesis. This includes an overview of the analysis chain used by all searches considered in Chapters 4 and 5 to turn a collection of protons at CERN into exclusion limits on supersymmetric

scenarios. Additionally, an overview of missing transverse momentum in ATLAS is provided, setting the stage for the machine-learning-based improvements in Chapter 6.

Concluding remarks are made in Chapter 7.

Chapter 2.

Theoretical background

2.1. The Standard Model

This thesis is concerned with the detection and discovery of fundamental particles. But what are particles? In the language of quantum field theory (QFT), they are field representations of the Poincaré group of *spacetime* symmetries (translations, rotations and Lorentz boosts) that carry internal angular momentum called *spin*. This includes scalars $\phi(x)$ with zero spin, Dirac spinors $\psi^a(x)$ with half-integer spin, and vectors $A_\mu(x)$ with integer spin. These fields are functions of space-time, and a ‘particle’ is a quanta of its corresponding field.

The Standard Model (SM) is currently the most successful theory of fundamental particles and their interactions. The SM is described by a Lagrangian which is invariant under local $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge transformations of the fields. Each of these gauge symmetries corresponds to a (vector) gauge field and (by Noether’s theorem [7]) a conserved charge. The $SU(3)_C$ gauge symmetry describes the strong interaction (QCD) and its generators correspond to 8 massless *gluons* which carry colour charge (red, green or blue). The $SU(2)_L$ gauge symmetry corresponds to three massless $W^{1,2,3}$ bosons which carry weak isospin, while $U(1)_Y$ corresponds to the massless B boson that carries weak hypercharge. Together, $SU(2)_L \times U(1)_Y$ describes the electroweak interaction [8–10].

The Standard Model Lagrangian is built from terms involving the fields associated with each particle. These terms must satisfy global Poincaré symmetry, gauge invariance, renormalisability and energetic stability (that there should be a vacuum state with the lowest possible energy). There are seven non-trivial classes of terms satisfying these

conditions which are combined to build the following Standard Model Lagrangian:

$$\mathcal{L}_{\text{SM}} = -\frac{1}{4}F^2 + \theta F * F + i\bar{\psi}\not{D}\psi + y\phi\bar{\psi}\psi + |D\phi|^2 + \mu^2\phi^2 + \lambda\phi^4 + h.c. \quad (2.1)$$

where sums over Lorentz, spinor, gauge group and fermion generation indices are suppressed for readability [11].

The Yang-Mills term $-\frac{1}{4}F^2$ is a kinetic term which governs the self-interactions of the gauge fields. Here F is the field strength tensor, $-igF = [D, D]$ where D is the covariant derivative, $D = \partial - igA$. There is a Yang-Mills term for each gauge group, with coupling g and gauge field A .

The term $\theta F * F$, where $*$ denotes a totally antisymmetric Lorentz contraction and θ is a dimensionless parameter, violates charge conjugation parity (CP) symmetry. However, constraints from measurements of the neutron electric dipole moment require $\theta < 10^{-10}$. This ‘fine-tuning’ is known as the strong CP problem [12]. This term has no consequences in perturbative theories so is not otherwise relevant to the physics in this thesis.

The Dirac term $i\bar{\psi}\not{D}\psi$ is the kinetic term for the Dirac spinors $\psi(x)$, describing their interactions with the gauge fields. Hence, this term governs the interactions of matter with the electromagnetic, weak and strong forces. Here the Feynman slash notation $\not{D} = \gamma^\mu \cdot D_\mu$ is used, where γ^μ are the gamma matrices which satisfy $\{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu,\nu}$ and $\eta^{\mu,\nu}$ is the Minkowski spacetime metric tensor.

Mass terms for the Dirac spinors and gauge fields are forbidden in the Standard Model Lagrangian by gauge invariance. Instead, the Higgs mechanism produces mass terms for these fields via their couplings to the Higgs field, ϕ . The Higgs field is a complex scalar doublet which is invariant under $SU(2)_L$ transformations. The kinetic term $|D\phi|^2$ describes the interaction between the Higgs field and the gauge fields, and the Higgs potential is $V(\phi) = \mu^2\phi^2 + \lambda\phi^4$, where $\mu^2 < 0$ and $\lambda > 0$. This potential has an infinite number of minima at a non-zero value of $\langle\phi\rangle = v = \frac{|\mu|}{\sqrt{\lambda}}$. Fixing the minimum, for example by arbitrarily choosing

$$\langle\phi\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}, \quad (2.2)$$

spontaneously breaks the $SU(2)_L$ symmetry. Now ϕ can be expanded around this minimum as

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix} \quad (2.3)$$

where h is the field associated with the Higgs boson. Evaluating $|D\phi|^2$ results in terms of the form $\frac{1}{2}m_Z^2 ZZ$ and $m_W^2 W^+ W^-$ where

$$W^\pm = \frac{1}{\sqrt{2}}(W^1 \mp iW^2) \quad \text{with mass } m = \frac{1}{2}vg_2 \quad (2.4)$$

and

$$Z = \frac{1}{\sqrt{g_1^2 + g_2^2}}(g_2 W^3 - g_1 B) \quad \text{with mass } m = \frac{1}{2}v\sqrt{g_1^2 + g_2^2}. \quad (2.5)$$

These combinations of the $W^{1,2,3}$ and B fields are identified as the $W^\pm$ and Z bosons. The fourth linearly independent combination is the photon,

$$A = \frac{1}{\sqrt{g_1^2 + g_2^2}}(g_2 W^3 + g_1 B) \quad \text{with mass } m = 0, \quad (2.6)$$

which is the generator of the unbroken $U(1)_Q$ symmetry, corresponding to the electromagnetic interaction. The familiar electric charge is a combination of the weak hypercharge and the third component of the weak isospin,

$$Q = I_3 + \frac{1}{2}Y. \quad (2.7)$$

To summarise the Higgs mechanism, the $SU(2)_L \times U(1)_Y$ symmetry is spontaneously broken to $U(1)_Q$ by fixing the vacuum expectation value of the Higgs field, generating masses for three of the four gauge bosons. Three of the four degrees of freedom associated with the Higgs doublet are used to give masses to the W^+ , W^- and Z bosons, and the remaining degree of freedom is the spin-0 Higgs boson. The photon and Z boson are mixtures of W^3 and B while the $W^\pm$ are mixtures of W^1 and W^2 . The W and Z bosons have a relatively large mass which makes the weak force very short-range. Additionally, the Yukawa interaction term in Equation 2.1, $y\phi\bar{\psi}\psi$, describes the interaction between the Higgs field and the Dirac spinors. After electroweak-symmetry breaking, this gives rise to mass terms for the Dirac spinors.

Name	Spin	Mass	Charge
Quarks			
u (up)	$\frac{1}{2}$	2.2 MeV	$+\frac{2}{3}$
d (down)	$\frac{1}{2}$	4.7 MeV	$-\frac{1}{3}$
c (charm)	$\frac{1}{2}$	1.3 GeV	$+\frac{2}{3}$
s (strange)	$\frac{1}{2}$	93 MeV	$-\frac{1}{3}$
t (top)	$\frac{1}{2}$	173 GeV	$+\frac{2}{3}$
b (bottom)	$\frac{1}{2}$	4.2 GeV	$-\frac{1}{3}$
Leptons			
e (electron)	$\frac{1}{2}$	0.511 MeV	-1
ν_e (e neutrino)	$\frac{1}{2}$	< 0.8 eV	0
μ (muon)	$\frac{1}{2}$	106 MeV	-1
ν_μ (μ neutrino)	$\frac{1}{2}$	< 0.19 MeV	0
τ (tauon)	$\frac{1}{2}$	1.78 GeV	-1
ν_τ (τ neutrino)	$\frac{1}{2}$	< 18.2 MeV	0
Gauge bosons			
γ (photon)	1	0	0
$W^\pm$	1	80.34 GeV	± 1
Z	1	91.19 GeV	0
g (gluon)	1	0	0
Scalar bosons			
h (Higgs)	0	125.25 GeV	0

Table 2.1.: The Standard Model particles and their properties. The anti-particles of the quarks and leptons (which have identical mass but opposite charge) are not shown. Mass values are taken from Reference [13].

The Standard Model particles are listed in Table 2.1, along with their spin, mass and electric charge. These particles can be divided into two categories: those with half-integer spin (*fermions*) and integer spin (*bosons*). The bosons include the aforementioned spin-0 Higgs boson and spin-1 gauge bosons. The fermions comprise matter particles and each has an *anti-matter* counterpart with the same properties except opposite sign electric charge. Fermions are represented by Dirac spinors in the Standard Model Lagrangian as described above.

The Standard Model fermions are in two chiral representations: left- or right-handed. The $SU(2)_L$ weak interaction only couples to left-handed fermions. In the chiral representation, the left-handed fermions are in a doublet while the right-handed fermions are singlets. The $W^\pm$ bosons couple exclusively to left-handed fermions, while the Z boson couples preferentially (but not exclusively) to left-handed fermions through its W^3 component (but also couples to charged right-handed fermions through its B component).

The fermions can be further subdivided into quarks and leptons which have different interactions. Leptons do not carry colour charge so do not interact via the strong force. The charged leptons (the electron, muon and tauon) interact via the electromagnetic and weak interactions while the neutral leptons, called *neutrinos*, only interact via the weak force and are massless. The six flavours of lepton can be split into three generations with identical properties other than mass. The left-handed leptons are arranged in three $SU(2)_L$ doublets (one for each generation) containing an electromagnetically charged and neutral lepton: $(\ell, \nu_\ell)_L$. The right-handed charged leptons on the other hand are singlets: ℓ_R . Right-handed neutrinos have not been observed, but their inclusion in the Standard Model with either Dirac or Majorana neutrino mass terms would explain the experimentally observed non-zero left-handed neutrino masses [14, 15].

The quarks carry both colour and electric charge and interact via the strong, weak and electromagnetic forces. There are six *flavours* which come in three generations of up- and down-type quarks with electric charge $+\frac{2}{3}$ and $-\frac{1}{3}$ respectively. The left-handed quarks are arranged in three $SU(2)_L$ doublets containing an up- and down-type quark: $(u, d)_L$. The right-handed quarks on the other hand are singlets: u_R, d_R . Quarks can never be observed in isolation due to *confinement*. This is because the QCD coupling constant α_S decreases as the momentum scale increases which gives QCD a property called *asymptotic freedom*. As a result, attempting to separate two coloured particles just creates potential energy from which a new quark-antiquark pair will be formed. Therefore, quarks produced in collider experiments result in a spray of hadrons called *jets*. Another consequence of asymptotic freedom is that QCD is non-perturbative at low energies, which means that numerical approximations are required to describe the structure of protons through parton-distribution functions (PDFs).

Additionally, the mass eigenstates of the fermions do not correspond to the weak interaction eigenstates, which leads to flavour mixing through the weak force. The dimensionless couplings y in the Yukawa terms parametrise the flavour mixing angles along with the fermion masses. For the quarks, this is described by the Cabbibo-Kobayashi-Maskawa (CKM) matrix [16, 17], parametrised by three mixing angles and a

complex phase. Similarly, the mixing between the neutrino flavour and mass eigenstates is encoded in the PMNS matrix [18, 19].

Overall there are twenty-six free parameters in the Standard Model that must be determined experimentally: the three gauge couplings g , two Higgs potential parameters μ and λ , twelve Yukawa couplings for fermion masses (six quarks, three charged leptons, three right-handed neutrinos), three angles and one phase of the CKM matrix, the four parameters defining the PMNS matrix, and the QCD θ term.

2.2. Problems with the Standard Model

All of the aforementioned particles have been observed¹, with the final piece being the Higgs boson discovery in 2012 [20, 21]. However, despite its overwhelming success at describing fundamental physics, there are several outstanding problems with the Standard Model. This includes ‘aesthetic’ problems such as the hierarchy problem and the strong CP problem which relate to the fine-tuning of free parameters. Additionally, there are several observations that the Standard Model fails to explain, including:

- Dark matter
- Insufficient CP-violation to explain matter-antimatter asymmetry
- The origin of neutrino masses
- Experimental tensions such as $g - 2$ [22, 23] and lepton flavour universality violation [24]
- Gravity is not included

All of these problems motivate new physics ‘beyond-the-Standard Model’ (bSM) which can be searched for at high-energy physics experiments. The two most relevant problems to the supersymmetric bSM scenarios discussed in this thesis are the hierarchy problem and dark matter, and further details of these will be provided in the following two sections.

¹Except for right-handed neutrinos.

2.2.1. Dark matter

Astrophysical phenomena across a range of energy scales motivate the existence of *dark matter* which is impacting galactic dynamics through the gravitational force but cannot be seen. At the galactic scale, this includes observed galactic rotation curves [25–27], gravitational lensing [28, 29] and galaxy cluster motion [30, 31]. At larger scales, dark matter is necessary for simulations of large-scale structure formation to be consistent with cosmological surveys [32].

It is difficult to solve these problems, which span multiple energy scales, with modifications to gravity alone. Instead, the dark matter hypothesis proposes a new particle (or multiple particles) which interacts through gravity but not the electromagnetic force. Measurements of the cosmic microwave background by the Planck Collaboration [33] show that 27% of the energy density in the universe is dark matter, while only 5% is the baryonic matter described by the Standard Model. The remaining 68% is thought to be *dark energy*, which drives the expansion of the universe.

One class of dark matter candidates are *weakly interacting massive particles* (WIMPs) which are stable (or very long-lived) and have no electromagnetic or strong interactions. Importantly, no Standard Model particle satisfies these properties except neutrinos. However, the structure formation in the early universe can only be explained by non-relativistic (‘cold’) dark matter, while neutrinos are relativistic (‘hot’). Hence, new WIMPs must be added to the Standard Model to fully explain dark matter.

In the thermal freeze-out mechanism, WIMPs were in thermal equilibrium with Standard Model particles during the early universe. At this time, the universe was very hot and the abundance of energy allowed for the production and annihilation of all particles. As the universe cooled, the heavier particles stopped being produced and decayed into lighter, stable particles, including WIMPs. A WIMP would not have decayed but could have continued to annihilate into Standard Model particles. As the universe expanded, the WIMPs in the universe spread out hence their annihilation rate decreased and eventually the dark matter density in the universe remained constant. This is known as *freeze-out*. The density of dark matter observed today is determined by the annihilation cross-section of the thermally-averaged self-annihilation $\chi + \chi \rightarrow \text{SM} + \text{SM}$ cross-section $\langle\sigma v\rangle$ [34]:

$$\Omega_\chi h^2 \approx \frac{0.1 \text{ pb}}{\langle\sigma v\rangle} \approx 0.1 \left(\frac{0.01}{\alpha}\right)^2 \left(\frac{m_\chi}{100 \text{ GeV}}\right)^2 \quad (2.8)$$

where m_χ is the WIMP mass and α the coupling. The observed dark matter relic abundance $\Omega_\chi h^2 \approx 0.12$ is obtained if WIMPs have weak scale masses and couplings.

This strongly motivates searches for WIMPs at the LHC, where there is enough centre-of-mass energy in the proton–proton collisions to produce WIMPs with weak-scale mass. Additionally, direct detection experiments search for WIMPs through their weak interactions with protons and neutrons, and indirect detection experiments look for an excess in cosmic rays from annihilating dark matter particles.

2.2.2. The hierarchy problem

The square of the Higgs boson mass receives radiative corrections from loop diagrams such as those in Figure 2.1. The size of the corrections, Δm_h^2 , for fermion and boson loops such as those shown in Figures 2.1a and 2.1b respectively, satisfy

$$\Delta m_{h,f}^2(\mu) = -\frac{|\lambda_f|^2}{8\pi^2} \Lambda_{UV}^2 + \dots, \quad \Delta m_{h,b}^2(\mu) = +\frac{|\lambda_b|^2}{16\pi^2} \Lambda_{UV}^2 + \dots \quad (2.9)$$

where λ_f and λ_b are the Yukawa coupling constants of the fermion/boson and Λ_{UV} an ultraviolet momentum cutoff used to regulate the loop integral [35]. Contributions such as these exist for every particle with a Yukawa coupling to the Higgs field. If new particles don't couple to the Higgs, there would still be 2-loop contributions through couplings to Standard Model particles which themselves couple to the Higgs.

Therefore, the Higgs mass will receive corrections from any energy scale containing new physics. At a minimum, new physics is expected at the Planck scale ($\mathcal{O}(10^{19})$ GeV), where quantum gravitational effects (which aren't described by the Standard Model) become dominant. This would correspond to corrections to the (squared) Higgs mass




Figure 2.1.: Fermion and scalar loop contributions to the Higgs mass.

of $\mathcal{O}(10^{38})$ GeV. Perfect cancellations of various loop contributions across 38 orders of magnitude would be required to result in the observed Higgs mass of 125 GeV. Such *fine-tuning* is considered very unlikely (or ‘unnatural’) and motivates an alternative explanation. To prevent issues with fine-tuning, there must be new physics at a scale of around a few TeV.

Note that the hierarchy problem is not a problem with the Standard Model itself, where the Higgs mass is simply an experimentally measured input parameter. Rather, any future bSM theory which allows the Higgs boson mass to be calculated will suffer from this problem and must include a mechanism to protect the Higgs mass from these large corrections.

The relative sign difference between the scalar and fermion loop contributions in Equation 2.9 provides one possible solution to the hierarchy problem. If there exists a symmetry between bosons and fermions, such that every fermion has an equal-mass boson partner (and vice versa), then all of the radiative corrections would cancel each other out. This proposed symmetry between bosons and fermions is called *supersymmetry*, and provides the basis for one of the most popular bSM theories. The rest of this chapter will outline the key theoretical and phenomenological features of supersymmetry, following References [35] and [36].

2.3. Supersymmetry

Supersymmetry (SUSY) [37–42] is a proposed extension to the Standard Model which introduces an additional space-time symmetry between bosons and fermions. For this symmetry to be satisfied, every particle must have a *superpartner* with the same charge, weak isospin, colour and mass, but spin differing by one-half. The only pair of Standard Model particles satisfying these requirements are the Higgs boson and the neutrino. However, these have different lepton numbers so, assuming lepton number is conserved, these can’t be superpartners. Therefore, supersymmetry predicts a new superpartner particle (‘sparticle’) for every SM particle, with spin differing by one-half. The radiative corrections to the Higgs mass from each fermion/boson are then cancelled by the corrections from each partner boson/fermion, solving the hierarchy problem [43–46].

2.3.1. The Minimal Supersymmetric Standard Model

The Minimal Supersymmetric Standard Model (MSSM) realizes supersymmetry while introducing the minimum number of new particles that results in a phenomenologically viable theory. The field content of the MSSM is summarised in Table 2.2 and consists of supermultiplets of SM fields plus their superpartners.

The SM leptons are partnered by ‘sleptons’ (selectrons, smuons, staus and their corresponding sneutrinos). These are arranged in chiral supermultiplets – there are both right- and left-handed sleptons for each of the right- and left-handed SM leptons. Similarly, the quarks are partnered by ‘squarks’ which are colour triplets and also in chiral supermultiplets.

The vector bosons in the SM reside in gauge multiplets with fermionic superpartners called *gauginos*. Each state in the SM gluon colour octet has a spin 1/2 partner called a gluino ($\tilde{g}$). The SM gauge fields associated with the $SU(2)_L \times U(1)_Y$ symmetry of the electroweak sector also have spin 1/2 partners: the neutral bino is the partner of the

Supermultiplet	Spin 0	Spin 1/2	Spin 1
	Higgs	Higgsinos	
H_u	(H_u^+, H_u^0)	$(\tilde{H}_u^+, \tilde{H}_u^0)$	–
H_d	(H_d^+, H_d^0)	$(\tilde{H}_d^+, \tilde{H}_d^0)$	–
	Squarks $\tilde{q}$	Quarks q	
Q	$(\tilde{u}, \tilde{d})_L$	$(u, d)_L$	–
U	$\tilde{u}_R$	u_R	–
D	$\tilde{d}_R$	d_R	–
	Sleptons $\tilde{\ell}$	Leptons ℓ	
L	$(\tilde{\nu}_e, \tilde{e})_L$	$(\nu_e, e)_L$	–
E	$\tilde{e}_R$	e_R	–
		Gaugino	Gauge boson
$\mathcal{G}$	–	$\tilde{g}^a$ gluino	g^a
$\mathcal{W}$	–	$\tilde{W}^a$ wino	W^a
$\mathcal{B}$	–	$\tilde{B}$ bino	B

Table 2.2.: Field content of the MSSM. Note that only the first of the three quark/squark and lepton/slepton generations are shown for clarity.

Names	Gauge eigenstates	Mass eigenstates	Spin
Higgs bosons	$H_u^0 \ H_d^0 \ H_u^+ \ H_d^-$	$h \ H \ A \ H^\pm$	0
Squarks	$\tilde{u}_L \ \tilde{u}_R \ \tilde{d}_L \ \tilde{d}_R$	(same)	0
	$\tilde{s}_L \ \tilde{s}_R \ \tilde{c}_L \ \tilde{c}_R$	(same)	0
	$\tilde{t}_L \ \tilde{t}_R \ \tilde{b}_L \ \tilde{b}_R$	$\tilde{t}_1 \ \tilde{t}_2 \ \tilde{b}_1 \ \tilde{b}_2$	0
Sleptons	$\tilde{e}_L \ \tilde{e}_R \ \tilde{\nu}_e$	(same)	0
	$\tilde{\mu}_L \ \tilde{\mu}_R \ \tilde{\nu}_\mu$	(same)	0
	$\tilde{\tau}_L \ \tilde{\tau}_R \ \tilde{\nu}_\tau$	$\tilde{\tau}_1 \ \tilde{\tau}_2 \ \tilde{\nu}_\tau$	0
Neutralinos	$\tilde{B}^0 \ \tilde{W}^0 \ \tilde{H}_u^0 \ \tilde{H}_d^0$	$\tilde{\chi}_1^0 \ \tilde{\chi}_2^0 \ \tilde{\chi}_3^0 \ \tilde{\chi}_4^0$	1/2
Charginos	$\tilde{W}^\pm \ \tilde{H}_u^\pm \ \tilde{H}_d^\pm$	$\tilde{\chi}_1^\pm \ \tilde{\chi}_2^\pm$	1/2
Gluino	$\tilde{g}$	(same)	1/2

Table 2.3.: The particles beyond-the-Standard Model (except h) introduced by the MSSM. The pre-electroweak symmetry breaking gauge eigenstates and post-electroweak symmetry breaking mass eigenstates are shown. Mixing within the first two generations is assumed to be negligible.

$U(1)$ gauge field, and the winos (a charged pair plus a neutral particle) are partners of the $SU(2)_L$ gauge fields. Additionally, two separate Higgs doublets are required to avoid gauge anomalies. Each of these doublets has a spin 1/2 superpartner called a ‘higgsino’.

The superpartner fields introduced so far are not necessarily the mass eigenstates. The binos, winos and higgsinos mix to form two charged states called *charginos*, denoted $\tilde{\chi}_{1,2}^\pm$, and four neutral states called *neutralinos*, denoted $\tilde{\chi}_i^0$ for $i = 1, 2, 3, 4$. The specific mixing determines the mass spectrum of the charginos and neutralinos.

The second scalar Higgs doublet results in additional Higgs boson states after electroweak symmetry breaking. Overall there are two CP-even neutral scalars h and H

(one of which is the Standard Model Higgs boson), one CP-odd neutral pseudoscalar A , and a charged scalar $H^\pm$.

The sleptons and squarks can mix with other sparticles of the same electric charge. However, most of the mixing angles are expected to be very small. The first- and second-generation squarks and sleptons are expected to be in 7 nearly mass-degenerate unmixed pairs: $(\tilde{e}_R, \tilde{\mu}_R)$, $(\tilde{\nu}_e, \tilde{\nu}_\mu)$, $(\tilde{e}_L, \tilde{\mu}_L)$, $(\tilde{u}_R, \tilde{c}_R)$, $(\tilde{d}_R, \tilde{s}_R)$, $(\tilde{u}_L, \tilde{c}_L)$, $(\tilde{d}_L, \tilde{s}_L)$. The third generation left- and right-handed stops, sbottoms and staus are expected to mix to give mass eigenstates $\tilde{t}_1$, $\tilde{t}_2$, $\tilde{b}_1$, $\tilde{b}_2$ and $\tilde{\tau}_1$, $\tilde{\tau}_2$. The gluino does not have suitable quantum numbers to mix with any other particle. The gauge and mass eigenstates of the undiscovered MSSM particles are summarised in Table 2.3.

The MSSM by default allows baryon and lepton number violating terms, which is inconsistent with the high degree of precision to which lepton and baryon number conservation have been confirmed. For example, baryon number violation could allow proton decay $p \rightarrow \pi^0 e^+$, which is inconsistent with the current experimental lower bound on the proton lifetime of 10^{34} years [47, 48]. To remedy this, a symmetry called R -parity [49, 50] is introduced, with quantum numbers assigned by

$$R \equiv (-1)^{3(B-L)+2S} = \begin{cases} +1 & , \text{ SM particles} \\ -1 & , \text{ SUSY particles} \end{cases} \quad (2.10)$$

where B is baryon number, L is lepton number and S is spin. All SM particles have an R -parity value of 1, while their superpartners have an R -parity of -1. R -parity values combine multiplicatively, therefore R -parity conservation implies that all sparticles must be produced in pairs from collisions of SM particles at collider experiments. R -parity conservation also requires that the lightest sparticle (LSP) be stable. The LSP is usually assumed to be the lightest neutralino. As a stable, neutral and weakly interacting particle, the neutralino is therefore a viable dark matter candidate. There are also supersymmetric scenarios that violate R -parity and use alternative solutions to the problem of proton-decay. However, all scenarios considered in this thesis conserve R -parity.

Since supersymmetry requires that each sparticle has the same mass as its SM partner, they should have been discovered simultaneously with their SM partners. However, no sparticles have yet been discovered, therefore SUSY must be a broken symmetry such that the superpartners can have different masses. Broken supersymmetry must still provide a solution to the hierarchy problem, so the breaking mechanism must be ‘soft’ such that the superpartner masses are not too dissimilar to their SM counterparts. The

mechanism for SUSY-breaking is unknown, though there are several competing theories. For example, mSUGRA [51–53] and GMSB [54–56] break SUSY via gravitational and gauge interactions respectively. The MSSM does not specify the mechanism for SUSY breaking and instead includes soft supersymmetry breaking terms which are characterized by over 100 parameters.

2.3.2. Electroweak supersymmetry at the LHC

The electroweak sector of supersymmetry includes the electroweak gauginos (the neutralinos and charginos – referred to collectively as ‘electroweakinos’) and sleptons. The coloured sparticles (squarks and gluinos) have strong couplings and therefore significantly larger production cross-sections than slepton or electroweakino production. However, if the gluino and squark masses are much heavier than low-mass electroweakinos, then SUSY production at the LHC would be dominated by direct electroweakino and slepton production.

Indeed, current ATLAS limits on squark and gluino production extend beyond the TeV scale [57–59], while the limits on electroweakinos are much weaker. Furthermore, the lightest electroweakinos are also expected to be the lightest sparticles. The LSP by definition is the lightest neutralino, which will be accompanied by an additional nearby neutralino and/or chargino if it has a significant wino or higgsino component. Additionally, dark matter constraints favour scenarios with an LSP mass at or below the TeV scale with additional charginos, neutralinos and/or sleptons close in mass to enhance co-annihilation mechanisms. This makes electroweak production of sparticles a promising and important probe of SUSY at the LHC and motivates it as the focus of this thesis. A brief summary of the key collider signatures of electroweak supersymmetry is outlined in the following (note that gluinos and squarks are not considered here).

As discussed above, R -parity conservation implies that sparticles are produced in pairs at collider experiments. Additionally, a two-body sparticle decay must always include one SM particle and one sparticle to satisfy R -parity conservation. Each neutralino and chargino has weak couplings to (slepton, lepton) pairs from its bino and wino field components, plus a gaugino-higgsino-Higgs coupling and $SU(2)_L$ gaugino-gaugino-vector boson coupling. Therefore, the most common two-body decays are:

$$\tilde{\chi}_i^0 \rightarrow Z\tilde{\chi}_j^0, h\tilde{\chi}_j^0, W\tilde{\chi}_j^\pm, \ell\tilde{\ell}, \nu\tilde{\nu} \quad (2.11)$$

$$\tilde{\chi}_i^\pm \rightarrow Z\tilde{\chi}_1^\pm, h\tilde{\chi}_1^\pm, W\tilde{\chi}_j^0, \ell\tilde{\nu}, \nu\tilde{\ell} \quad (2.12)$$

Note that other decays are also possible via the additional Higgs bosons A , H and $H^\pm$.

In cases where all of these two-body decays are kinematically forbidden due to small mass splittings, the three-body decays

$$\tilde{\chi}_i^0 \rightarrow ff\tilde{\chi}_j^0, \quad \tilde{\chi}_i^0 \rightarrow ff'\tilde{\chi}_k^\pm, \quad \tilde{\chi}_i^\pm \rightarrow ff'\tilde{\chi}_j^0, \quad \tilde{\chi}_2^\pm \rightarrow ff\tilde{\chi}_1^\pm \quad (2.13)$$

occur through offshell gauge bosons, where f is a quark or lepton.

The sleptons decay into their lepton partner plus a chargino or neutralino:

$$\tilde{\ell} \rightarrow \ell\tilde{\chi}_i^0, \quad \tilde{\ell} \rightarrow \nu\tilde{\chi}_i^\pm, \quad \tilde{\nu} \rightarrow \nu\tilde{\chi}_i^0, \quad \tilde{\nu} \rightarrow \ell\tilde{\chi}_i^\pm. \quad (2.14)$$

Therefore, after the production of a pair of sparticles, a chain of decays will occur which produce h , W and Z bosons, plus leptons and quarks, and end in two LSPs². The presence of the LSP results in missing transverse momentum (p_T^{miss}), a collider observable which will be introduced in Section 3.4.6, as it is assumed to be neutral and stable. Since h , W and Z bosons also decay into leptons and quarks, the general signature is n leptons + m hadronic jets + p_T^{miss} . Standard Model processes involving W and Z bosons provide the most important backgrounds for these signatures, as W and Z decays can produce both leptons/jets and neutrinos as a source of p_T^{miss} .

2.3.3. Simplified models

The MSSM contains 120 free parameters, which is too many to be scanned exhaustively and compared to data. Instead, collider searches are usually optimized and interpreted using models that make assumptions to reduce these parameters. Many searches for sparticles use ‘simplified models’ which assume that only a small number of SUSY particles (usually two) are kinematically accessible and all others decoupled. Branching ratios are typically set to either 0% or 100% in simplified models to target specific decay chains. Additionally, electroweakinos are usually assumed to be a pure bino, wino or higgsino state. These assumptions allow the generation of events for SUSY signal models in a 2D grid (with the mass of a SUSY particle on each axis). The final ‘result’ of a

²Alternative signatures can also arise for scenarios involving long-lived particles that involve a sparticle traversing or leaving a detector before decaying.




Figure 2.2.: The observed and expected exclusion limits reported by an ATLAS search for pair-production of charginos with a leptonic final state [60], interpreted in a simplified model.

search (assuming no statistically significant signal is observed) is then a contour in this 2D plane containing all models which can be excluded.

For example, Figure 2.2 shows the exclusion contour for an ATLAS search for pair-production of charginos [60]. This simplified model assumes that only the lightest chargino $\tilde{\chi}_1^\pm$ and neutralino $\tilde{\chi}_1^0$ are kinematically accessible and includes just the $\tilde{\chi}_1^+ \tilde{\chi}_1^- \rightarrow W^+ \tilde{\chi}_1^0 W^- \tilde{\chi}_1^0$ production process and decay. The chargino is assumed to be purely wino and the neutralino purely bino. A statistical combination of searches for this simplified model is performed in Chapter 5.

The simplified model approach is very useful for designing targeted searches for specific sparticles, and the resulting 2D exclusion contours provide an intuitive interpretation of the results. However, it neglects much of the more complicated phenomenology of the MSSM. For example, relaxing the simplified model assumptions made when producing the limits in Figure 2.2 would reveal less than 100% exclusion for scenarios with chargino and neutralino masses within the reported contour. Therefore, it is important to interpret search results in more general SUSY models to give a complete picture of the current constraints on SUSY parameter space.

2.3.4. The phenomenological MSSM

The phenomenological MSSM (pMSSM) [61] is a middle ground between simplified models and the full 120-parameter MSSM. The pMSSM makes the following assumptions, motivated by experimental constraints and general features of possible SUSY breaking mechanisms, to reduce the MSSM parameters to 19:

- 1) No new sources of CP-violation** This corresponds to ensuring the soft SUSY breaking parameters in the MSSM are real. Sources of CP-violation from beyond the CKM matrix are tightly constrained by experimental limits on the electron and neutron electric dipole moments [62, 63] and in the K system. Note that this restriction means that the pMSSM alone cannot explain baryogenesis which does require some additional CP-violation beyond the Standard Model CKM matrix.
- 2) No flavour-changing neutral currents** Flavour-changing neutral currents (FCNCs) are forbidden at tree-level in the Standard Model and are tightly constrained by experimental results. In the pMSSM, Minimal Flavour Violation [64] is imposed by making all sfermion mass matrices and trilinear couplings diagonal.
- 3) First and second generation universality** Experimental data (for example from $K^0 - \bar{K}^0$ mixing), severely constrains the mass splitting between the first and second generation sfermions. In the pMSSM, it is assumed that the soft-SUSY breaking scalar masses are the same for the first and second generations and the trilinear couplings are zero. Note there are no such constraints on the third generation masses or trilinear couplings.
- 4) R -parity conservation** As discussed above, this prevents problems related to baryon and lepton number violation.
- 5) The LSP is the lightest neutralino** This ensures that the lightest neutralino is a viable dark matter candidate. This is not always required of pMSSM scenarios but is imposed for the study in Chapter 4.

The 19 parameters of the pMSSM are summarised in Table 2.4. The pMSSM has enough free parameters to capture some of the richer phenomenology present in the MSSM that is absent in simplified models, such as more complex decay chains. It also has a moderate enough number of parameters to be used to interpret search results.

Parameter	Description
$m_{\tilde{L}_1} (=m_{\tilde{L}_2})$	Left-handed slepton (first two gens.) mass
$m_{\tilde{e}_1} (=m_{\tilde{e}_2})$	Right-handed slepton (first two gens.) mass
$m_{\tilde{L}_3}$	Left-handed stau doublet mass
$m_{\tilde{e}_3}$	Right-handed stau mass
$m_{\tilde{Q}_1} (=m_{\tilde{Q}_2})$	Left-handed squark (first two gens.) mass
$m_{\tilde{u}_1} (=m_{\tilde{u}_2})$	Right-handed up-type squark (first two gens.) mass
$m_{\tilde{d}_1} (=m_{\tilde{d}_2})$	Right-handed down-type squark (first two gens.) mass
$m_{\tilde{Q}_3}$	Left-handed squark (third gen.) mass
$m_{\tilde{u}_3}$	Right-handed top squark mass
$m_{\tilde{d}_3}$	Right-handed bottom squark mass
M_1	Bino mass parameter
M_2	Wino mass parameter
μ	Higgs-Higgsino mass parameter
M_3	Gluino mass parameter
A_t	Trilinear top coupling
A_b	Trilinear bottom coupling
A_τ	Trilinear τ lepton coupling
M_A	Pseudoscalar Higgs boson mass
$\tan \beta$	Ratio of the vacuum expectation values of the two Higgs doublets

Table 2.4.: Summary of the 19 parameters of the pMSSM.

The nature of the electroweakinos in the pMSSM is dictated by the bino, wino and higgsino mass parameters. If these parameters are well separated, the neutralinos and charginos will be split into a bino-like singlet of mass approximately $|M_1|$ containing a single neutralino, a wino-like triplet of two oppositely charged charginos and a neutralino with mass approximately $|M_2|$, and a higgsino-like quartet of two charginos and two neutralinos with mass approximately $|\mu|$. Here ‘bino-like’ is defined as having a bino fraction greater than 50%, and similar for ‘wino-like’ and ‘higgsino-like’. Simplified models typically assume well-separated and pure wino/bino/higgsino neutralinos and charginos. The pMSSM allows the electroweakinos to be a mixture if the mass parameters are not separated.

In Chapter 4, a selection of ATLAS Run 2 SUSY searches (which were originally interpreted using simplified models) will be reinterpreted to constrain the electroweak sector of the pMSSM.

Chapter 3.

Experimental background: From protons to p -values

“If your experiment needs statistics, you ought to have done a better experiment”

— E. Rutherford

A theory of fundamental physics, as with any scientific model, is only as good as its agreement with data. To test the veracity of the Standard Model, and proposed extensions such as supersymmetry, predictions must be compared with data from high energy particle collisions such as the proton–proton collisions produced at the Large Hadron Collider (LHC). But how can a mathematically abstract Lagrangian representing a particle physics model be tested with the complex and messy products of a proton–proton collision? The process begins by accelerating and colliding protons at the LHC. Then, the particles produced in the collisions are detected and reconstructed by a particle detector such as ATLAS. Simulated data is used as a theoretical prediction and a statistical analysis is performed to determine whether the data provides evidence of new physics. This chapter outlines the steps in this procedure and techniques used for an ATLAS analysis, which underpins all the searches which are reinterpreted in Chapters 4 and 5.

3.1. The Large Hadron Collider

The LHC is the largest and highest energy particle collider in the world. It consists of a 27 km tunnel lying beneath the Swiss-French border at a mean depth of 100m. The LHC primarily collides protons (but occasionally heavy ions such as lead) which originate from a simple bottle of hydrogen.

Hydrogen atoms are ionised by an electric field to produce protons, which are then passed through a series of particle accelerators. First, LINAC2 accelerates the protons to 50 MeV, followed by the Proton Synchrotron Booster, Proton Synchrotron and Super Proton Synchrotron which take the protons to energies of 1.4 GeV, 25 GeV and 450 GeV respectively. Then, the 450 GeV protons are injected into the LHC ring and accelerated to their final collision energy. This series of accelerators are displayed in Figure 3.1, along with the rest of the CERN particle accelerator complex.

Two proton beams travel in opposite directions around the circular LHC ring, steered by dipole bending magnets. Each beam consists of bunches of $\mathcal{O}(10^{11})$ protons at 25 ns intervals. Quadrupole focusing magnets focus the beams at four *interaction points* where they cross, leading to proton–proton collisions.

At each interaction point lies a particle detector designed to detect and measure the products of the proton collisions. ATLAS [66] and CMS [67] are general-purpose detectors used to make precision measurements of Standard Model quantities and search for physics beyond-the-Standard Model. LHCb [68] is the only forward-backward asymmetric detector and its physics programme focuses on heavy-flavour hadrons. Finally, ALICE [69] is used to study the quark-gluon plasma produced in heavy ion collisions.

To observe rare production processes, such as those predicted by supersymmetry, a large number of collision events are required. The total number of collision events delivered by the LHC is given by $N = L\sigma$ where σ is the interaction cross-section and L is the integrated luminosity: $L = \int \mathcal{L}_{\text{inst.}} dt$. Here $\mathcal{L}_{\text{inst.}}$ is the instantaneous luminosity, which reaches $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ at the LHC. The integrated luminosity is often referred to simply as luminosity.

During the first data-taking run of the LHC between 2008–2012 (referred to as *Run 1*), ATLAS recorded a luminosity of 4.57 fb^{-1} and 20.3 fb^{-1} for pp collisions at $\sqrt{s} = 7 \text{ TeV}$ and $\sqrt{s} = 8 \text{ TeV}$ respectively. The analyses considered in this thesis use data from the 2015–2018 period (referred to as *Run 2*), where the collision energy was increased to




Figure 3.1.: The particle accelerator complex at CERN [65].

$\sqrt{s} = 13 \text{ TeV}$ and 139 fb^{-1} integrated luminosity was delivered – the luminosity for the Run 2 data-taking period is displayed in Figure 3.2a. At the time of writing, the third data-taking run (*Run 3*) has just begun with $\sqrt{s} = 13.6 \text{ TeV}$ collisions.

A proton is composed of a number of quarks and gluons (collectively called *partons*) and its momentum is split between these constituents. A collision between two protons can involve a *hard scatter*, which is a parton-parton interaction with large momentum transfer, plus underlying *soft* (lower momentum transfer) interactions between the remaining partons. Particles of interest (such as W bosons, Z bosons, Higgs bosons or SUSY particles) have high mass and hence are only produced in interactions with high momentum transfer - i.e. the hard scatter. Since the longitudinal dynamics of the




Figure 3.2.: Displayed are: (a) Total integrated luminosity delivered by the LHC (green), recorded by ATLAS (yellow) and good for physics analysis (blue) during Run 2; (b) the luminosity weighted distribution of the mean number of interactions per bunch crossing for each year of data-taking during Run 2. Reproduced from Reference [70].

partons involved in the hard scatter are not known, variables are used which characterise particle dynamics in the transverse plane, such as transverse momentum. Typically the z -axis is defined as the direction of the incoming protons, such that the transverse momentum is $p_T = \sqrt{p_x^2 + p_y^2}$.

In addition to interesting hard scatter events, the data collected from the LHC is contaminated with proton–proton collisions known as *pile-up*. Soft interactions between other protons in the same bunch crossing give rise to additional soft particles superimposed on the final state of the hard scatter. This *in-time* pile-up is parametrised by the number of reconstructed primary collision vertices, N_{PV} , in the bunch crossing. ATLAS is also sensitive to the remnants of previous bunch crossings as the detector readout rate is lower than the bunch-crossing frequency. This *out-of-time* pile-up is parametrised by the mean number of interactions per bunch crossing, $\langle\mu\rangle$. Distributions of $\langle\mu\rangle$ for Run 2 are shown in Figure 3.2b. Pile-up interactions increase with instantaneous luminosity and, in general, produce lower momentum particles that contribute to background. Various experimental techniques are applied to reduce pile-up contamination in physics analyses, some of which will be described in Section 3.4.

3.2. The ATLAS detector

The ATLAS (A Toroidal LHC ApparatuS) experiment [66] at the LHC is a multi-purpose particle detector used to identify particles produced in the proton–proton collisions and measure their energy and momentum. Its original goals were to discover the Higgs boson and search for new physics beyond-the-Standard Model. The former was achieved in 2012 [20].

The ATLAS detector uses a cylindrical coordinate system with its origin at the centre of the detector, z -axis along the beam pipe, x -axis pointing from the centre of the detector to the centre of the LHC ring, and y -axis pointing upwards. Cylindrical coordinates (r, ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the z -axis. The pseudorapidity is defined in terms of the polar angle θ as $\eta = -\ln \tan \frac{\theta}{2}$. ATLAS is hermetic in ϕ and has forward-backward coverage up to $|\eta| < 4.9$, allowing for the reconstruction of nearly every energetic particle in an event.




Figure 3.3.: Computer-generated image of the ATLAS detector, including each subdetector. Image from Reference [66].

The subcomponents of ATLAS are arranged in concentric layers which take care of different measurements and particle identifications. A cutaway image of ATLAS with each component labelled is shown in Figure 3.3. The inner detector measures the charge and momentum of a charged particle from the direction and radius of its curved path through a magnetic field provided by the superconducting solenoid magnet. The electromagnetic calorimeter measures the energies of electrons and photons by absorbing them. Similarly, the hadronic calorimeter measures the energies of charged and neutral hadrons like pions. The Muon Spectrometer is the outermost layer and provides a stand-alone measurement of the momenta of muons with a dedicated magnetic field from the toroid magnets. Additionally, three forward detectors (LUCID, the CDZ and ALFA) measure the luminosity and detect forward neutrons from heavy ion collisions. Further details on each of these subcomponents of ATLAS will be provided in the following sections.

3.2.1. Inner detector

The inner detector (ID) [71, 72], shown in Figure 3.4, consists of a tracking system in a 2 T axial magnetic field provided by a superconducting solenoid [73], covering the pseudorapidity range $|\eta| < 2.5$. A charged particle traverses a curved path (referred to as a *track*) through the magnetic field, leaving behind hits in the sensors of several subsystems. The particle's charge and momentum can be determined from the curvature of the track. The ID consists of three subdetectors: the high-resolution pixel tracker, the strip-based semiconductor tracker (SCT), and the transition radiation tracker (TRT).

Pixel detector The silicon pixel detector [74] is the closest ATLAS component to the interaction point and allows track reconstruction at a high particle rate. It contains 1700 modules, each of which consists of a 2×6 cm sensor with 47,000 pixels. These are arranged in three concentric barrel layers, providing at least three hits per traversing charged particle. Two endcaps containing three disks each cover the forward regions. The insertable B-layer [75] was added before Run 2 and provides an extra hit near the beam pipe for improved tracking resolution, primary vertex reconstruction, secondary vertex finding and flavour tagging.

Semiconductor tracker (SCT) The SCT [76] contains six million silicon strips across 4088 modules arranged in four concentric barrel layers plus two endcaps with nine disks each. It provides precise three-dimensional spatial measurements for charged




Figure 3.4.: Schematic diagram of the ATLAS inner detector. Image from Reference [66].

particle tracks and plays an essential role in track reconstruction, charged particle identification and b -tagging.

Transition radiation tracker (TRT) The TRT [77] consists of 350,000 straw drift tubes with 4 mm diameter and covers the $|\eta| < 2$ region. These are oriented parallel to the beam pipe in the barrel region and perpendicular to the beam pipe on the endcaps. The straws are filled with a mixture of xenon, carbon dioxide and oxygen gas. A particle passing through ionises the gas to produce a measurable signal. Electron identification is possible from the detection of X-ray photons from transition-radiation.

3.2.2. Calorimeters

Calorimeters measure the energy of particles by absorbing them and converting their energy into measurable charge or scintillation light via ionisation or excitation. Unlike the trackers used in the ID, calorimeters can detect neutral as well as charged particles. ATLAS uses sampling calorimetry, where alternating layers of passive and active material absorb particles and detect the resulting signals respectively. This allows for segmentation and compactness at the expense of energy resolution. Dense materials are used in the passive layers to instigate a cascade of secondary particles. The energy measured is proportional to the number of particles in the resulting shower which are detected by the




Figure 3.5.: Schematic diagram of the ATLAS calorimeter system. Image from Reference [66].

active medium. The ATLAS calorimetry system [78] is shown in Figure 3.5. It covers $|\eta| < 4.9$ and extends radially to 4.25m from the beam-pipe.

The electromagnetic calorimeter detects electromagnetically interacting particles, such as photons and electrons. The barrel covers $|\eta| < 1.475$ and consists of 1024 pairs of lead absorbers arranged in an accordion geometry. Cavities are filled with the active material, Liquid Argon (LAr) [79], for fine granularity energy measurements of photons and electrons. The endcaps cover $1.375 < |\eta| < 3.2$. For photons, the dominant energy loss mechanism (for energies above a few MeV) is pair production $\gamma \rightarrow e^+e^-$. On the other hand, high energy electrons lose energy primarily through Bremsstrahlung (photon emission). Therefore, a photon or electron entering the electromagnetic calorimeter can be identified from a characteristic cascade of e^+e^- pairs and photons through several iterations of these processes, called an *electromagnetic shower*. The original particle can further be identified as a photon or electron by looking for matched tracks in the ID – an electron will have left tracks while a neutral photon will not.

Charged and neutral hadrons that make it through the electromagnetic calorimeter reach the hadronic calorimeter, where they produce a *hadronic shower* from various interaction processes with the calorimeter material. The TILE calorimeter [80] covers the outer region and uses steel absorbers with plastic scintillating tiles arranged in three layers. Photons produced from the interactions of charged particles in the scintillator are collected by wavelength-shifting optical fibres at the end of cells. Photomultiplier

tubes at the end of each row of cells measure these signals. The deposited energy is proportional to the number of detected photons. The barrel covers the region $|\eta| < 1.0$ while the extended barrel covers $0.8 < |\eta| < 1.7$.

The hadronic end-cap and forward calorimeter [81] use LAr rather than plastic scintillators due to its radiation hardness. The hadronic end-cap (HEC) covers $1.5 < |\eta| < 3.2$ and has two disks per end cap with a copper absorber. The forward calorimeters cover the forward region $3.1 < |\eta| < 4.9$ and have three layers. The first layer uses a copper absorber for sensitivity to electromagnetic showers and tungsten for the outer two layers for better containment of hadronic showers.

3.2.3. Muon spectrometer

The muon spectrometer [82], shown in Figure 3.6, is a large tracking system designed to make precise momentum measurements of muons which pass through the electromagnetic and hadronic calorimeters. ATLAS makes a stand-alone measurement of the muon momentum (independent of the ID), which requires a magnetic field in the muon region to produce curved muon tracks. This is provided by three 0.5–1 T toroidal magnets situated in the barrel and end-cap regions.




Figure 3.6.: Schematic diagram of the ATLAS muon spectrometer. Image from Reference [66].

Precise measurements are made by monitored drift tubes (MDTs) and cathode strip chambers (CSCs). MDTs are 30 mm diameter tubes filled with argon and carbon dioxide gases. Traversing muons ionise the gas to produce electrons which drift towards the central anode wire to provide a readout. There are three concentric MDT layers in the barrel region and the endcap consists of four wheels placed orthogonal to the beam axis to give a total coverage up to $|\eta| < 2.7$. The MDT timing resolution (maximal drift time) is 700 ns. In the region $2.0 < |\eta| < 2.7$, where the most muons are expected, cathode strip chambers (CSCs) are used to provide an improved timing resolution of 40 ns.

Overall, the muon momentum resolution from the precision tracking chambers is around 10% up to a few TeV, and muons with energies as low as 3 GeV can be measured. However, this precision comes at the cost of slow read-out. Therefore a coarser but faster measurement is required to achieve an adequate read-out rate for triggering. This is provided by resistive plate chambers (RPCs) in the barrel region. RPCs are pairs of 2 mm-spaced parallel resistive plates. The gaps are filled with gas that is ionised by passing muons to form avalanches of electrons which are detected. Thin gap chambers (TGCs) cover the end-cap region to extend the coverage up to $|\eta| < 2.4$ and consist of multi-wire proportional chambers. The RPCs and TGCs have a response time of 1.5 ns and 4 ns respectively which is ideal for triggering and well below the 25 ns bunch spacing to enable a robust association of a given measurement with a particular bunch crossing.

3.2.4. Forward detectors

Three smaller sets of detectors provide coverage in very forward regions:

The Luminosity measurement using Cherenkov Integrating Detector (LUCID) is 17m from the interaction point. Its main purpose is to detect inelastic pp scattering in the forward direction for measuring the instantaneous and integrated luminosity and for monitoring beam conditions on-line. It consists of 20 aluminium tubes filled with gas – as particles traverse the gas Cherenkov light is emitted and collected by a photomultiplier to produce a signal.

The zero degree calorimeter (ZDC) is located 140 m from the interaction point and detects very forward neutrons ($|\eta| > 8.3$) from heavy-ion collisions.

Absolute Luminosity measurement For ATLAS (ALFA) also measures the luminosity in ATLAS by measuring the elastic scattering rate in the forward region. It consists of a scintillating fibre tracker and is located 240 m from the interaction point.

3.3. Triggering

The LHC produces proton bunch crossings at a rate of 40 MHz. Saving all of these events presents a monumental data storage challenge. To tackle this, the most interesting events are identified and saved using a two-level triggering system. In general, interesting events are those containing high- p_T objects (electrons, muons, photons, taus or jets) or high missing transverse momentum.

First, the level 1 (L1) hardware-based trigger uses low-level detector information from the calorimeters and muon spectrometer to make a coarse but quick ($< 2.5 \mu\text{s}$) initial trigger decision and reduce the event rate to 100 kHz. The electronics for the L1 trigger are located in the ATLAS cavern to reduce the latency in the trigger decision. Then, the software-based high-level trigger (HLT) reduces the event rate to 1 kHz. The HLT has enough time (around 300 ms per event) to perform a fast reconstruction of signals from a wider array of detector subsystems for use in the triggering decision.

3.4. Object reconstruction

The detector components described thus far provide a raw dataset of experimental signals such as hit counts and energy deposits. But as particle physicists, we are interested in the kinematics of *particles*. The procedure to obtain momenta of particles and other physics ‘objects’ from raw detector data is called *object reconstruction*. This section briefly describes the reconstruction of the key objects in a typical ATLAS analysis: electrons, muons, photons, taus, hadronic jets, and missing transverse momentum (p_T^{miss}).

3.4.1. Low-level objects

Clusters, tracks and primary vertices are the low-level objects used to build each of the high-level objects that follow. Clusters are groups of calorimeter cells with a significant signal that are considered to originate from the same physics object. Details on the clustering used for electron/photon reconstruction (‘egamma’ clusters) and hadronic jet reconstruction (‘topoclusters’) will be given in the relevant following sections. Tracks are formed from raw hits left in the pixel, SCT and/or TRT (and muon spectrometer for muon tracks) from the traversal of a charged particle with $p_T > 0.5 \text{ GeV}$ and $|\eta| < 2.5$. Vertices are identified from the origin of multiple tracks. The primary vertex of an event

is the vertex with the highest Σp_T^2 from all tracks with $p_T > 500$ MeV associated to it – this is assumed to be the origin of the hard scatter.

Each of the following high-level objects are reconstructed entirely independently. Therefore, it's possible for the same calorimeter cluster or track to be identified with multiple reconstructed objects (for example to both an electron and a jet), leading to double counting. An analysis-dependent overlap removal procedure is usually applied to address such ambiguities.

3.4.2. Electrons and photons

Electrons and photons are reconstructed in parallel from a similar set of detector information [83–85]. Deposits in the electromagnetic calorimeter are used to form ‘egamma’ clusters. These are constructed with a ‘sliding window’ algorithm using fixed $\Delta\eta \times \Delta\phi$ windows. The $\eta \times \phi$ space of the EM calorimeter is divided into a 200×256 grid of cell ‘towers’ spanning the three calorimeter layers. A window spans 3×5 towers, and windows with a total $E_T > 2.5$ GeV are used as seeds from which clusters are formed by recursively adding neighbouring windows with localised energy deposits.

Egamma clusters are then matched to tracks from the ID. This is achieved by extrapolating the tracks from their last measured point in the ID to the middle of the LAr calorimeter and evaluating a distance metric with nearby clusters. An unmatched cluster is assumed to be a photon (which doesn't leave hits in the ID). A hard scatter photon which has converted into an e^+e^- pair is identified by looking for two close-by oppositely-charged tracks originating from a secondary vertex displaced from the primary vertex.

Electron isolation and identification requirements are imposed to reduce contamination from hadronic decays (‘non-prompt’ leptons), photon conversions and hadrons misidentified as leptons. Loose, medium and tight isolation working points are defined based on the ratio of the electron's energy to the rest of the cluster and track energies in a surrounding cone-shaped region. Identification algorithms use a likelihood-based discriminant based on the shower shape, degree of leakage into the hadronic calorimeters and other cluster and track properties.

To ensure that the electron track comes from the primary vertex, analyses often impose that the transverse impact parameter significance ($|d_0|/\sigma(d_0)$) is less than five and the longitudinal impact parameter (Δz_0) satisfies $|\Delta z_0 \sin \theta| < 0.5$ mm. Here Δz_0 is

the difference between the z coordinates of the primary vertex and the point of closest approach between it and the track.

3.4.3. Muons

Muons are reconstructed from track segments in the muon spectrometer. For tracks with $|\eta| < 2.5$, these are additionally matched to, and combined with, tracks in the ID for a more precise measurement and identification than using the muon spectrometer alone. [86, 87]

Similar to electrons, identification and isolation working points are defined to discriminate prompt muons from those originating from hadronic decays. The requirements for muons are typically $|\Delta z_0 \sin \theta| < 0.5$ mm and $|d_0|/\sigma(d_0) < 3$.

3.4.4. Hadronic jets

Due to QCD colour confinement, quarks and gluons produced in the proton–proton collisions are not observed directly but hadronise, producing a collimated spray of pions and other hadrons known collectively as a *jet*.

Jets are reconstructed primarily from the deposits of these hadrons in the electromagnetic and hadronic calorimeters. These deposits form topological clusters (topoclusters) [88] consisting of 3D groupings of topologically connected calorimeter cells. Note that these are distinct from the 2D-grouped ‘egamma’ clusters used for electron/photon reconstruction.

Topoclusters are formed from a ‘seed’ cell with $|E_{\text{cell}}| > 4\sigma_{\text{noise}}$ where E_{cell} is the cell signal and σ_{noise} is the average expected noise in the cell. Neighbouring cells are then added to the cluster if they satisfy $|E_{\text{cell}}| > 2\sigma_{\text{noise}}$. The energy of each topocluster is the sum of the energy of all constituent cells, accounting for potential overlap between nearby topoclusters. The angular position (η, ϕ) of the topocluster is calculated from the energy-weighted barycentre of the cell positions.

Topoclusters are calibrated to produce an accurate energy measurement. The exact calibration procedure applied depends on whether the cluster is electromagnetic- or hadronic-like, as the deposits from each interact differently with the calorimeters. Topoclusters are initially calibrated to the electromagnetic energy response scale. An

additional hadronic calibration is then applied to clusters classified as hadronic in origin as the calorimeter response to hadrons is lower than that for electrons.

Topoclusters have good energy resolution but poor spatial resolution as they include no tracking information. This led to the development of a second type of jet input object: *particle-flow* (PFlow) [89]. Tracks are matched to topoclusters using *ghost association* to combine the information from both the tracking detectors and calorimeters. For tracks with $p_T < 100$ GeV, the track energy is added and the expected calorimeter energy deposition of the track is calculated and subtracted from its matched topocluster to avoid double-counting. For $p_T > 100$ GeV the tracking performance is degraded so the track energy is not used and the subtraction is disabled. Unmatched topoclusters are assumed to originate from neutral particles, which don't leave any tracks. Charged PFlow objects (PFOs) which are not matched to the primary vertex are assumed to be pile-up and are removed. PFlow jets, therefore, have improved pile-up rejection compared to topocluster jets.

Clustering algorithms are used to combine the low-level jet inputs into a single jet object. In ATLAS, jets are predominantly reconstructed using the anti- k_t jet clustering algorithm [90]. This takes an input parameter R , which loosely corresponds to the 'width' of the jet in the (η, ϕ) plane. Typical jets are reconstructed with $R = 0.4$. However, heavier particles like top quarks and W , Z , and Higgs bosons can produce high- p_T boosted jets with wide angles. These 'large-radius' jets are clustered with $R = 1.0$ [91] in some analyses. The four-vector of the final jet is calculated as the vector sum of the four-vectors of the final jet constituents.

Jets originating from a b -quark (referred to as b -jets) are of particular interest as the dominant Higgs boson decay mode is via a pair of b -quarks and top quarks almost always decay via a b -quark. A b -quark has a typical lifetime of around 1×10^{-12} s which is longer than the lighter quarks, and it travels approximately 3 mm before decaying. This produces a jet which contains a secondary vertex from the displaced b -quark decay. This can be used to distinguish a b -jet from jets originating from a light-quark or gluon, via a procedure known as b -tagging.

Jet reconstruction is strongly affected by pile-up contamination since pile-up mainly consists of low-energy hadronic activity which can overlap with calorimeter signals from a hard scatter jet. To reduce this, various *pile-up mitigation algorithms* are available that attempt to subtract pile-up contributions from the constituent jet inputs. Additionally, a Jet Vertex Tagger (JVT) [92] algorithm uses tracking information to quantify the

likelihood that a given jet originates from pile-up- (JVT = 0) or the hard scatter (JVT = 1). Requirements on the JVT of a jet can improve the rejection of pile-up jets. Similarly, a ‘forward Jet Vertex Tagger’ (fJVT) discriminates hard scatter and pile-up jets in the forward region ($|\eta| > 2.5$).

3.4.5. Tau leptons

Similar to b -quarks, taus are long-lived enough to travel from their production point before decaying (producing a slightly displaced vertex) but short-lived enough to still decay within the detector volume. A tau decaying leptonically ($\tau \rightarrow \mu\nu_\mu\nu_\tau$ or $\tau \rightarrow e\nu_e\nu_\tau$) can’t be distinguished from prompt electrons or muons as the neutrinos aren’t detected. Only hadronically decaying taus (about 65% of all tau decays) can be identified.

Hadronic taus are reconstructed as an $R = 0.4$ jet. A machine-learning-based classifier is then used to discriminate jets originating from a tau from hadronic jets using tracking and vertex input information [93].

3.4.6. Missing transverse momentum

Missing transverse momentum (p_T^{miss}) represents the total transverse momentum of all undetected particles produced in proton–proton collisions [94]. In many beyond-the-Standard Model searches, including all the SUSY searches outlined in Chapters 4 and 5, p_T^{miss} is associated with the production of dark matter candidates (such as the lightest neutralino in the pMSSM) – since dark matter does not interact via the electromagnetic or strong force it would escape the detector without being detected. Similarly, Standard Model neutrinos are not directly detected by ATLAS, and instead show up as p_T^{miss} .

The p_T^{miss} is reconstructed using the fact that the total transverse momentum of the protons prior to colliding is zero, so by conservation of momentum the total transverse momentum of the products of the collision should also sum to zero. The p_T^{miss} is therefore equal to the magnitude of the negative vectorial sum of the transverse momentum of all the visible particles produced in the hard scatter. This quantity is a challenge to reconstruct, as it involves every particle that the ATLAS detector observes and all of the detector subsystems. Limitations due to the detector resolution and pile-up contributions result in an observation bias towards non-vanishing values of p_T^{miss} even for events without

any invisible particles. A non-zero reconstructed p_T^{miss} for an event with zero *true* p_T^{miss} ($p_T^{\text{miss, True}}$) is referred to as *fake* p_T^{miss} (as opposed to *real* p_T^{miss}).

The reconstructed p_T^{miss} consists of contributions from hard objects (fully reconstructed electrons, photons, tau leptons, muons and jets) plus soft signals (tracks associated with the hard scatter vertex but not with any of the hard objects). This is summarised in Equation 3.1.

$$\mathbf{p}_T^{\text{miss}} = - \sum_{\text{electrons}} \mathbf{p}_T^e - \sum_{\text{photons}} \mathbf{p}_T^\gamma - \sum_{\tau\text{-leptons}} \mathbf{p}_T^{\tau_{\text{had}}} - \sum_{\text{muons}} \mathbf{p}_T^\mu - \sum_{\text{jets}} \mathbf{p}_T^{\text{jet}} - \sum_{\text{unused tracks}} \mathbf{p}_T^{\text{track}} \quad (3.1)$$

Note that $\mathbf{p}_T^{\text{miss}}$ here is a 2-vector with x and y components. Many searches make use of the components themselves, in addition to the magnitude p_T^{miss} , as they are used to define various other variables such as the transverse mass (m_T) and stransverse mass (M_{T2}) [95, 96]. The use of separate terms in Equation 3.1 allows the p_T^{miss} calculation to benefit from the better momentum resolution of photons and electrons relative to jets. However, each type of hard object has its own independent reconstruction procedure. This can lead to double counting in the p_T^{miss} calculation if certain calorimeter/tracking signals are used in the reconstruction of more than one object. To ensure p_T^{miss} is reconstructed from mutually exclusive detector signals, the hard objects are reconstructed in a predefined order, with an algorithm to reject objects associated with detector signals that have already been used. The order of priority for the hard objects is usually as follows: electrons, photons, τ_{had} and finally jets. Muons are not included in this procedure as there is little overlap between muons and other objects since they are reconstructed using ID and MS tracks rather than calorimeter signals¹. Each analysis will define its own selection criteria for the reconstruction of these objects. Lower priority objects are rejected if they share signals with a higher priority object that has already been reconstructed. Additionally, there is a more refined strategy to scale $\mathbf{p}_T^{\text{jet}}$ in the case of partial overlap of calorimeter signals associated with jets and other objects [94].

Finally, all tracks which remain and are not associated with a hard object are included in a soft term (the final term of Equation 3.1). Note that the soft term can also be built from calorimeter signals as opposed to tracks, which has the advantage of including neutral particles but at the expense of reduced pile-up suppression.

¹ There is a separate overlap-removal procedure for muons which overlap with jets.

The hard and soft components of the p_T^{miss} are impacted by pile-up contributions, which leads to fake p_T^{miss} . This is primarily due to pile-up contamination in the jet term. Several jet selection working points have been developed to deal with pile-up, and each consists of a different requirement on the jets entering the p_T^{miss} calculation. There is a trade-off here: stricter requirements on the jets entering the jet term result in improved pile-up rejection, however this also risks rejecting hard-scatter jets. The tracks associated with jets which fail the requirements instead enter the track-based soft term, but any neutral objects contributing to these jets will be completely ignored in the p_T^{miss} calculation. Therefore, the optimal jet selection to use for a given event depends on its topology. For example, for a topology like $t\bar{t}$ which produces many jets it is optimal to use a looser jet selection when reconstructing p_T^{miss} to ensure all hard-scatter jets are included in the p_T^{miss} calculation, even at the expense of some pile-up contamination. For a topology like $Z \rightarrow \mu\mu$ on the other hand, which contains fewer jets, a stricter jet selection ensures improved rejection of pile-up jets without rejecting jets from the hard scatter [3]. The definitions of a selection of the working points available to ATLAS analyses (of Run 2 data) are listed in Table 3.1. In Chapter 6, a machine-learning-based combination of several working points is presented.

Table 3.1.: Jet selections for ATLAS p_T^{miss} reconstruction working points.

Working point	Selections			fJVT for jets with $2.5 < \eta < 4.5$ and $p_T < 120$ GeV
	p_T [GeV] for jets with: $ \eta < 2.4$	p_T [GeV] for jets with: $2.4 < \eta < 4.5$	JVT for jets with $ \eta < 2.4$	
<i>Loose</i>	> 20	> 20	> 0.5 for $p_T < 60$ GeV jets	-
<i>Tight</i>	> 20	> 30	> 0.5 for $p_T < 60$ GeV jets	< 0.4
<i>Tighter</i>	> 20	> 35	> 0.5 for $p_T < 60$ GeV jets	-
<i>Tenacious</i>	> 20	> 35	> 0.91 for $20 < p_T < 40$ GeV jets > 0.59 for $40 < p_T < 60$ GeV jets > 0.11 for $60 < p_T < 120$ GeV jets	< 0.5

Many searches for SUSY and other bSM scenarios that employ p_T^{miss} additionally make use of p_T^{miss} *significance* to separate signal processes with real p_T^{miss} from backgrounds which produce fake p_T^{miss} . The p_T^{miss} significance is defined by taking the ratio of p_T^{miss} and a measure of its resolution. The current state-of-the-art ATLAS p_T^{miss} significance

variable is an object-based $p_{\text{T}}^{\text{miss}}$ significance, defined in Reference [97] as

$$\mathcal{S} = \frac{|p_{\text{T}}^{\text{miss}}|^2}{\sigma_L^2(1 - \rho_{LT}^2)}. \quad (3.2)$$

Here σ_L^2 is the momentum resolution of objects in the direction longitudinal to the $\mathbf{p}_{\text{T}}^{\text{miss}}$ 2-vector and ρ_{LT}^2 is the correlation factor for object measurements longitudinal and transverse to the $\mathbf{p}_{\text{T}}^{\text{miss}}$ 2-vector. This variable therefore quantifies the likelihood that a given $p_{\text{T}}^{\text{miss}}$ is arising due to genuine invisible particles as opposed to resolution effects in the objects in the $p_{\text{T}}^{\text{miss}}$ sum.

3.5. Analysing ATLAS data

The experimental hardware and software-based object reconstruction procedure outlined so far produces data containing the kinematics of particles produced in proton–proton collisions. But how is this used to test theoretical models of particle physics, such as supersymmetry? In this section, a brief overview will be given of the analysis methods used in searches for new physics, including all the searches reinterpreted in Chapters 4 and 5. This includes the generation of simulated data to act as a theoretical prediction for comparison to data, and the statistical techniques used to test the Standard Model against bSM signals.

3.5.1. Simulated data

Simulated data is used as a theoretical prediction to be compared with data. The production of simulated data is handled by the ATLAS simulation framework [98].

First, events are generated by simulating a specified hard-scatter process to produce final state primary particles. This includes simulation of the interactions between the quarks and gluons in the colliding protons and the subsequent parton showering, hadronisation and decays into stable particles. This is performed with Monte-Carlo (MC) generators like PYTHIA [99] and SHERPA [100]. At this stage, the events are at ‘particle’-level (also referred to as ‘truth’-level), as the impact of detector effects is absent.

Next, the impact of the ATLAS detector is simulated by passing generated events through a GEANT4 simulation [101] of the ATLAS detector to produce ‘hits’ – a prediction

of how each particle in such an event would have traversed the detector and deposited its energy in the various components. Digitisation then subjects GEANT4 hits to the detector response to provide a prediction of the detector component readouts that would have been observed for such an event. This results in simulated data in the same format as the raw data from real events. The effects of pile-up are also replicated by overlaying simulated pile-up events on simulated data at this stage. From this point, simulated data is reconstructed in the same way as real data, as outlined in Section 3.4.

A fast simulation of ATLAS is often performed using ATLFastII [102, 103]. This uses the full GEANT4 simulation for the ID and MS only and replaces the calorimeter simulation with FASTCALOSIM [104] which smears the truth-level information using resolutions measured in studies of fully simulated events. This significantly reduces the CPU time required to simulate the large signal grids required in bSM searches.

3.5.2. Event selection

Most searches for SUSY (and other bSM scenarios) involve defining several disjoint bins in kinematic observables, and then counting the number of events observed in each bin compared to MC-based (or data-driven) signal and background predictions. To observe typically small SUSY signals above large Standard Model backgrounds in the data, kinematic regions must be identified where the signal is more likely to be found (and background less likely). These ‘signal regions’ are obtained by optimising selections on kinematic variables which reject background while accepting signal events based on studies with simulated data. Key variables used in the SUSY searches considered in this thesis include p_T^{miss} , transverse mass (m_T) and stransverse mass (M_{T2}) [95, 96].

Due to imperfections in theoretical modelling and the detector simulation, simulated data does not always perfectly reflect the distributions produced by real data. To address this, ‘control’ regions are defined which contain selections enriched in particular background processes in a kinematic phase space that is similar to the signal regions. Observations in a control region can then be used in the statistical fit to calculate an appropriate *normalisation factor* to scale the background process to match the data. This scaling factor is then applied to the background in the signal region. Validation regions are also defined to ensure this correction works as expected and the observed data matches the fitted backgrounds with the normalisation factors found.

3.5.3. Statistical analysis

The complex data processing steps outlined so far are impacted by randomness at every step. For example, proton–proton collision events are inherently probabilistic and classical randomness impacts the detector response. Additionally, theoretical calculations and detector modelling are imperfect. Therefore, how can a bona fide discovery be distinguished from a statistical fluctuation or mis-modelling in part of the data-processing chain?

To answer this question, a statistical model of the data is required that reflects these uncertainties,

$$f(\mathbf{x} \mid \boldsymbol{\phi}) = f(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\theta}). \quad (3.3)$$

This is the probability distribution function (PDF) of the event rates $\mathbf{x}$ in each analysis region given model parameters $\boldsymbol{\phi}$. The observed data is one value $\mathbf{x}_{\text{obs}}$ of $\mathbf{x}$ drawn from the PDF. The model parameters include both *parameters of interest* (POIs) $\boldsymbol{\mu}$ and nuisance parameters (NPs) $\boldsymbol{\theta}$. POIs are parameters we are interested in measuring, such as the *signal strength*: ‘how much’ signal is present in the data. NPs are assumed to be random variables that affect the event rates but with unknown true values, such as the background normalisation factors introduced previously and systematic uncertainties. Systematic uncertainties quantify our incomplete knowledge of the statistical model. For example, there are systematic uncertainties related to jet energy resolution, and theoretical uncertainties quantifying the uncertainty in the cross-section of processes included in the analysis.

Building a statistical model

The analyses considered in this thesis involve the simultaneous measurement of disjoint binned distributions (channels) observed as event counts n . For each channel, the expected number of observed events is the sum over a number of processes (samples) including signal and background. The event count in a single bin obeys a Poisson distribution of the form

$$\text{Pois}(n; \nu) = \frac{\nu^n e^{-\nu}}{n!} \quad (3.4)$$

where n is the count and ν is the expectation value.

SUSY searches are typically more complex than a single bin counting experiment, involving the measurement of multiple bins simultaneously. The distribution of the counts in multiple bins is a product of Poisson distributions. The expectation value for the count in each bin is based on MC-based or data-driven predictions for each sample, but is also subject to modifiers that cause a deviation in the expected event rates and quantify what we don't know about the data through the model parameters ϕ . The event rate for bin b of channel c can therefore be defined as

$$\nu_{cb}(\phi) = \sum_{s \in \text{samples}} \nu_{scb}(\boldsymbol{\eta}, \boldsymbol{\chi}) = \sum_{s \in \text{samples}} \left(\prod_i \kappa_{scb}^i(\boldsymbol{\eta}, \boldsymbol{\chi}) \right) (\nu_{scb}^0(\boldsymbol{\eta}, \boldsymbol{\chi}) + \sum_j \Delta_{scb}^j(\boldsymbol{\eta}, \boldsymbol{\chi})). \quad (3.5)$$

Here the total rates are given by a sum over individual sample rates which are determined by a nominal rate ν_{scb}^0 and a set of multiplicative and additive rate modifiers κ^i and Δ^j . The nominal rates ν_{scb}^0 are determined from Monte-Carlo simulation and/or data-driven methods. In Equation 3.5 the parameters ϕ have been partitioned into *free parameters* $\boldsymbol{\eta}$ and *constrained parameters* $\boldsymbol{\chi}$. POIs are free parameters to be determined from a fit to the data. Nuisance parameters can be free, in which case they are the normalisation factors introduced in Section 3.5.3, or constrained. Constrained NPs represent the systematic uncertainties in the model. The degree to which constrained parameters can impact the expected event rates is controlled through constraint terms $c_\chi(a_\chi|\chi)$ in the PDF. Here a_χ is an *auxiliary measurement*, typically from a control region. Therefore the total observation is $\boldsymbol{x} = (\boldsymbol{n}, \boldsymbol{a})$ where $\boldsymbol{n}$ is the channel data and $\boldsymbol{a}$ the auxiliary measurements.

It is usually assumed that $\boldsymbol{n}$ and $\boldsymbol{a}$ are independent, such that their joint PDF is the product of their individual PDFs. Putting everything together, the PDF for a multi-binned analysis takes the form [105]

$$f(\boldsymbol{x} | \phi) = f(\boldsymbol{n}, \boldsymbol{a} | \boldsymbol{\eta}, \boldsymbol{\chi}) = \prod_{c \in \text{channels}} \prod_{b \in \text{bins}} \text{Pois}(n_{cb} | \nu_{cb}(\boldsymbol{\eta}, \boldsymbol{\chi})) \prod_{\chi \in \boldsymbol{\chi}} c_\chi(a_\chi | \chi) \quad (3.6)$$

$\nu_{cb}(\boldsymbol{\mu}, \boldsymbol{\chi})$ are the expected rate of events as a function of the model parameters, specified in Equation 3.5. The constraint terms control the impact of constrained parameters on the event rates, and are usually parametrised by a Gaussian $G(\hat{\chi}, \chi_0, \sigma_\chi)$ with central value $\chi_0 = 0$ and uncertainty σ_χ .

Maximum likelihood estimation

The parameters of this statistical model are determined through *maximum likelihood estimation* (MLE). The *likelihood* function is constructed by fixing the event counts in the statistical model to the observed data,

$$L(\phi) = f(\mathbf{x} = \mathbf{x}_{\text{obs}} \mid \phi). \quad (3.7)$$

The parameters of the model ϕ are then determined by finding the values $\hat{\phi}$ that maximise the likelihood. $\hat{\phi} = \arg \max L(\phi)$ is a function of the data and there's no guarantee that this procedure obtains the 'true' parameter values (as there may be fluctuations in the data), but it does provide a sensible estimate.

To simplify this calculation, usually

$$\lambda = -2 \log L(\phi) \quad (3.8)$$

is minimised rather than maximizing $L(\phi)$ directly. Since the likelihood typically has a Gaussian form, taking the negative log converts a product of exponentials into a sum of parabolic terms which simplifies the optimisation task. A non-linear minimization is performed to find $\hat{\phi}$ by solving

$$\frac{\partial \lambda}{\partial \phi_i} = 0. \quad (3.9)$$

For typical ATLAS analyses, including the statistical fits performed in Chapters 5 and 4, computational packages such as Minuit [106] and SciPy [107] are used to perform this minimization.

After this procedure, the parameters have been measured and can be compared to their nominal values. In particular, since parameters associated with systematic uncertainties are constrained to obey a Gaussian distribution, their measured value should not deviate too far from zero. A pull is defined as $(\hat{\chi} - \chi_0)/\sigma_\chi$, where χ_0 is the pre-fit expected value of the nuisance parameter (typically zero) and $\hat{\chi}$ is the post-fit value. Pulls within 1σ are expected due to fluctuations in the data, but larger pulls than this can indicate mis-modelling or other issues with the fit.

Hypothesis testing

The statistical model can also be used to test particular theoretical models against the observed data through hypothesis testing. A hypothesis test takes a null hypothesis H_0 and alternative hypothesis H_1 and determines if H_0 is disfavoured or favoured by the data compared to H_1 . The hypotheses to be tested in a bSM search are based on specific values of the signal strength parameter μ in the statistical model. This represents the level of the signal (ie. the bSM model being searched for) present in the data. A value of $\mu = 1$ indicates that the amount of signal events predicted by MC simulation is present, while a value of $\mu = 0$ indicates that no signal is present and the data can be described by the Standard Model background processes alone. When testing hypotheses, rather than finding the value of μ most consistent with the data, the goal is to test whether a specific value of μ can be ruled out by the data.

If the observed data shows an excess over the background estimation, one can test for discovery of new physics. Here the null hypothesis is $\mu = 0$ (the Standard Model) and the alternative hypothesis is $\mu > 0$ (that there is signal present in the data). The central question is whether the observed excess is due to the presence of signal or due to a fluctuation in the background estimate. This is quantified using p -values, which represent the probability of finding data as or more incompatible with the given hypothesis.

The Neyman-Pearson Lemma states that when comparing two hypotheses H_0 and H_1 , the optimal discriminator is the likelihood ratio $L(H_0)/L(H_1)$ [108]. The negative log of the likelihood ratio is usually used as the test statistic [109] to simplify computation,

$$q_0 = -2 \log \left(\frac{L(H_0)}{L(H_1)} \right). \quad (3.10)$$

For a test in a bSM search of a specific value of μ that also includes the impact of nuisance parameters the profile log-likelihood ratio is used to build appropriate test statistics,

$$\lambda(\mu') = -2 \log \left(\frac{L(\mu = \mu', \hat{\boldsymbol{\theta}}(\mu'))}{L(\hat{\mu}, \hat{\boldsymbol{\theta}})} \right). \quad (3.11)$$

Here $\hat{\boldsymbol{\theta}}$ is the best-fit value of the nuisance parameters under the assumption $\mu = \mu'$ and $\hat{\mu}$ and $\hat{\boldsymbol{\theta}}$ the overall best-fit values (where μ is not fixed in the fit).

For a discovery test where the Standard Model hypothesis $\mu = 0$ is tested against the signal hypothesis $\mu > 0$, the profile likelihood is used to build the test statistic

$$q_0 = \begin{cases} \lambda(0) & , \quad \hat{\mu} \geq 0 \\ 0 & , \quad \hat{\mu} < 0 \end{cases} \quad (3.12)$$

This is defined such that the Standard Model hypothesis can't be rejected by a deficit in the data such that $\hat{\mu} < 0$.

The p -value is computed as

$$p_0 = \int_{q_0^{\text{obs}}}^{\infty} f(q_0 | \mu = 0) dq_0 \quad (3.13)$$

where $f(q_0 | \mu)$ is the PDF of the test statistic. The p -value quantifies how likely it is to obtain data as or more discrepant than the data observed, under the assumption that the Standard Model hypothesis is true.

In general, a p -value of 0.03 indicates ‘evidence’ of new physics. A discovery requires a p -value of 3×10^{-7} – a 1 in 3.5 million chance of the result being due to a statistical fluctuation. Such a high threshold is necessary to provide protection against a false discovery due to the look-elsewhere effect – searches typically cover multiple regions so the chance of having a fluctuation *somewhere* is higher – or mis-modelled systematic uncertainties [110].

The searches considered in this thesis did not observe significant excesses above the Standard Model background yield predictions. So instead of testing for discovery, the analyses excluded signal models by placing upper limits on the signal yields. Typically this involves reporting a 95% confidence limit (CL) on the signal strength. To allow for the possibility of the best-fit $\hat{\mu}$ being negative (for which there is a 50% chance in the absence of signal in the data), an alternative definition of the profile log-likelihood,

$$\tilde{\lambda}(\mu') = \begin{cases} -2 \log \frac{L(\mu=\mu', \hat{\theta}(\mu'))}{L(\hat{\mu}, \hat{\theta})} & , \quad \hat{\mu} \geq 0 \\ -2 \log \frac{L(\mu=\mu', \hat{\theta}(\mu'))}{L(\mu=0, \hat{\theta}(0))} & , \quad \hat{\mu} < 0 \end{cases} \quad (3.14)$$

is used to build the test statistic

$$q_{\mu_0} = \begin{cases} \tilde{\lambda}(\mu_0) & , \quad \mu_0 \geq \hat{\mu} \\ 0 & , \quad \mu_0 < \hat{\mu} \end{cases} . \quad (3.15)$$

The p -value for the agreement between the data and the tested value of $\mu = \mu_0$ is

$$p_{\mu_0} = \int_{q_{\mu_0}}^{\infty} f(q_{\mu}|\mu) dq_{\mu} . \quad (3.16)$$

For a given threshold α , a signal strength $\mu = \mu_0$ is excluded at confidence level $(1 - \alpha)$ if $p_{\mu_0} < \alpha$. For example, for a 95% confidence limit, α is set to 0.05. However, if the analysis has low sensitivity to a particular model, a downward fluctuation in the data can lead to false-exclusion when no sensitivity was expected. To combat this, a ‘modified p -value’, $\text{CL}_s = \frac{p_1}{p_0}$ is used [111]. Here p_0 is the p -value for the background-only hypothesis ($\mu = 0$) and p_1 the p -value for $\mu = 1$. A signal model is typically considered ‘excluded’ if the observed $\text{CL}_s < 0.05$.

In most SUSY searches, a CL_s value is obtained for a 2D grid of signal models in a simplified model parameter space. By interpolating the results between the grid points, a 2D contour can be drawn indicating the region of simplified model scenarios that the analysis is expected to be sensitive to (the expected limit) and that have been excluded (the observed limit). An example exclusion contour from Reference [60], seen previously in Chapter 2, is shown in Figure 3.7.

3.6. Reinterpretation tools

The statistical fits and model interpretation for most SUSY searches, including all of the analyses considered in Chapters 4 and 5 of this thesis, were performed using a set of simplified SUSY signal models. However, other signal hypotheses could in principle be tested against the background hypothesis produced by each of these analyses. For example, the limitations of simplified models compared to the pMSSM were outlined in Chapter 2, and in Chapter 4, a large number of pMSSM model points are evaluated with several SUSY searches to determine if they are excluded. Additionally, theory colleagues developing bSM models may wish to test their models with existing searches. To achieve this, several tools are available to ‘reinterpret’ ATLAS analyses with new




Figure 3.7.: Observed and expected exclusion limits on a SUSY simplified model of chargino-pair production, reproduced from Reference [60].

signal models. This requires information about the statistical model of the background processes used in the analysis, plus the ability to process new signal models with the analysis event selections to obtain signal yields in each kinematic region of the analysis. The reinterpretation tools used in this thesis are briefly introduced here.

3.6.1. SimpleAnalysis

SIMPLEANALYSIS [112] is an analysis framework that runs on truth-level event samples (ie. events that have not been run through ATLAS simulation). It includes the option to smear the four-vectors of each particle to approximate the impact of detector effects. This uses the UpgradePerformanceFunctions, which consist of object reconstruction and identification efficiencies and resolutions parameterised in η and ϕ .

This allows for a very fast approximation of analysis region yields for a given signal model, as the CPU-intensive detector simulation does not need to be performed. In general, this results in yields which agree with a full (or fast) detector simulation within about 20%. The publication of a SIMPLEANALYSIS implementation is a requirement for approval of all ATLAS SUSY searches and, unlike the full ATLAS simulation framework, the SIMPLEANALYSIS framework is available to non-ATLAS users.

3.6.2. JSON likelihoods

The likelihood is one of the most important products of an analysis as it contains all the information about the backgrounds, yields and uncertainties which are necessary to reproduce the analysis results. ATLAS now publishes full analysis likelihoods built using the HISTFACTORY template [105, 113] in JSON format. This file contains the background and observed yield estimates for every analysis region, along with systematic uncertainty deviations.

PyHF [114, 115] is designed to read these JSON files and use them to run statistical fits in Python. PyHF can be used to run statistical fits for new signal models by inserting (“patching”) expected signal yields (and systematic variations if desired) into the JSON file. For example, in Chapter 4 approximate signal yields for pMSSM points are obtained using SIMPLEANALYSIS, patched into a JSON likelihood for each analysis and fitted using PyHF to obtain an approximate CL_s value.

Additionally, ATLAS and CMS have recently introduced simplified likelihoods [116, 117]. This makes several simplifications, such as merging all backgrounds into a single sample and reducing all nuisance parameters in the full likelihood to a single constrained parameter. This significantly increases the speed of the fit, at the price of some accuracy.

3.6.3. RECAST

RECAST [118, 119] is a framework for re-running existing analyses with new signal models. Unlike SIMPLEANALYSIS, RECAST runs over event samples that have been passed through the ATLAS simulation chain to produce an accurate yield estimation for each region. A RECAST implementation of an analysis allows for a full re-execution of the event selection and statistical analysis for a new signal sample. This requires the preservation of the software environment and relevant commands for the full analysis workflow. This is achieved using Docker containers [120] which package a full software environment into a portable format (including the operating system, file system and the application itself). These docker containers, along with commands for running the relevant analysis code, are chained together into a workflow with user-defined inputs (typically the signal model event samples and cross-section). This completely defines the analysis pipeline, from event selection to final statistical fitting.

In Chapter 4, RECAST will be used to obtain an accurate CL_s value for pMSSM models where the approximate value obtained from SIMPLEANALYSIS yields patched into a JSON likelihood is considered insufficient.

3.7. Machine learning methods

A high energy physics experiment such as ATLAS involves the analysis of huge datasets, which makes it ripe for the application of machine learning (ML). In recent years, machine learning algorithms have been used to assist in almost every step of the ATLAS analysis chain outlined so far, including (but not limited to) improved object reconstruction and calibration [121], jet flavour tagging [122], and signal vs. background discrimination [123]. In Chapter 6, neural networks will be used to improve missing transverse momentum reconstruction performance. A comprehensive overview of the machine learning methods used in ATLAS is beyond the scope of this thesis. Instead, this section provides a brief introduction to neural networks and the regression and classification techniques used in Chapter 6.

3.7.1. Neural networks

The ML methods used in this thesis involve *supervised learning*, where the goal is to approximate a function $\mathbf{y} = f(\mathbf{x})$ where $\mathbf{y}$ is a ‘target’ and $\mathbf{x}$ is a set of input variables (often referred to as *features*). A neural network can be used to approximate the function $f(\mathbf{x})$ via a series of layers containing nodes in a structure summarised in Figure 3.8. The structure begins with an input layer with a node for each input feature and ends with an output layer that produces a prediction of the targets $\mathbf{y}$. Between the input and output layers are a series of *hidden layers* containing a number of nodes. Each line in Figure 3.8 has an associated weight w_{ij} which is used to perform a transformation of the values entering each node. The output at node $a_j^{(i)}$ is

$$a_j^{(i)} = \sigma \left(\sum_{k=1}^m w_{j,k} \cdot a_k^{(i-1)} + b_j^{(i-1)} \right), \quad (3.17)$$

where $a_k^{(i-1)}$ are the node values from the previous layer and $b_j^{(i-1)}$ a bias term from the previous layer to complete the linear transformation. Non-linearity is introduced by the activation function σ . The choice of activation function depends on the specific problem,




Figure 3.8.: The basic neural network structure of an input layer, hidden layers and an output layer.

but often the Rectified Linear Activation Unit (ReLU),

$$\text{relu}(x) = \begin{cases} x & , \quad x \geq 0 \\ 0 & , \quad x < 0 \end{cases}, \quad (3.18)$$

is recommended for hidden layers. It has been shown that a neural network with a single hidden layer can approximate any function arbitrarily closely given enough nodes [124]. In practice, stacking multiple hidden layers often reduces the number of neurons required to represent a given function.

To produce a close approximation of $\mathbf{y} = f(\mathbf{x})$ a neural network is *trained* with example data. During training, the weights $w_{i,j}$ are adjusted such that the outputs $\mathbf{y}$ better match the targets $\hat{\mathbf{y}}$. The metric used to quantify the agreement between $\mathbf{y}$ and $\hat{\mathbf{y}}$ is called the *loss function*,

$$\mathcal{L} = \mathcal{L}(\hat{\mathbf{y}}, f(\mathbf{x})), \quad (3.19)$$

where $\hat{\mathbf{y}}$ is the true value of the target, and $f(\mathbf{x})$ is the prediction. Various loss functions are available to suit different tasks.

Gradient descent methods are used to iteratively adjust the hidden layer weights to minimise the loss function. For iteration t this typically involves a transformation of the form

$$\mathbf{w}^{(t)} = \mathbf{w}^{(t-1)} - \eta \frac{\partial \mathcal{L}}{\partial \mathbf{w}} \quad (3.20)$$

where η is the learning rate. The derivative $\frac{\partial \mathcal{L}}{\partial \mathbf{w}}$ is calculated through a procedure known as backpropagation [125]. Various modifications to this basic algorithm exist for performing the optimization. The Adam optimizer [126] is typically recommended and is used in Chapter 6.

Hyperparameters are parameters that affect the training and performance of a neural network but are not learned during training. Instead, they are set to a fixed value. This includes the learning rate, the number of hidden layers and nodes, and the initial weight values. Hyperparameter tuning usually involves testing a variety of values and picking those which result in the best performance.

3.7.2. Regression

A regression problem involves predicting a continuous variable given a set of inputs. For example, one could predict truth-level quantities (such as jet energy) from detector-level quantities to calibrate the objects reconstructed by ATLAS.

A typical choice of loss function for a regression problem is the mean squared error (MSE),

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2. \quad (3.21)$$

The mean squared error will particularly punish a network for outlier predictions with large errors. This has the advantage of reducing the presence of such outliers, but the disadvantage of focusing the network on uncommon outliers at the expense of improved predictions in the more populated regions of the training set. Therefore in some cases, the mean absolute error (MAE) loss function is used instead,

$$\mathcal{L}_{\text{MAE}} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|. \quad (3.22)$$

As a middle ground, the Huber loss function [127] can be used. This is defined as

$$\mathcal{L}_{\text{Huber}} = \begin{cases} \frac{1}{2}(y - \hat{y})^2 & , \quad |y - \hat{y}| \leq \delta \\ \delta|y - \hat{y}| - \frac{1}{2}\delta^2 & , \quad \text{otherwise} \end{cases} \quad (3.23)$$

where δ is a hyperparameter which would typically be set to around $\delta = 1.5$.

3.7.3. Classification

Classification problems involve picking a category given a set of features. For example, a search for supersymmetry may use classifiers to classify events as ‘signal-like’ or ‘background-like’ to aid in signal region definitions. This is an example of a *binary* classification. A *multi-class* classifier has more than two possible classes.

The standard loss function choice for classification problems is the categorical cross-entropy loss, defined as

$$\mathcal{L}_{\text{CCE}} = \sum_C t_C \log(P_C) \quad (3.24)$$

where t_C is the truth-label for class C (1 if C is the true class, 0 otherwise) and P_C is the value output by the network for class C . The Softmax activation function,

$$\sigma(\mathbf{z})_i = \frac{e^{z_i}}{\sum_{j=1}^K e^{z_j}}, \quad (3.25)$$

is usually used for the output layer, where z_i is the i th output. Combined with the categorical cross-entropy loss, this imposes that the output scores sum to one and can be interpreted as the probability that the data point belongs to each class.

Chapter 4.

Interpretation of $\sqrt{s} = 13$ TeV ATLAS searches in the pMSSM

4.1. Introduction

The phenomenological MSSM (pMSSM) was introduced in Chapter 2. It is a phenomenologically viable model of SUSY which includes all sparticles from the MSSM but makes several experimentally and theoretically motivated assumptions to reduce the 120 parameters of the MSSM to 19.

ATLAS performed a range of searches for sparticles during the first two data-taking runs of the LHC. No statistically significant signals were observed and exclusion limits were placed on SUSY model parameters – typically the sparticle masses in simplified model scenarios. However, as discussed in Chapter 2, the simplified models used to produce these limits made assumptions which are likely not true if SUSY is realised in nature. In searches for electroweak production, this typically includes assuming that each electroweakino is a pure bino/wino/higgsino state with just one decay mode. Whilst exclusion limits using simplified models provide an easily interpretable picture of the sensitivity of analyses to specific areas of parameter space, they are far from an exhaustive exploration of the MSSM.

At the end of Run 1 a reinterpretation of the ATLAS SUSY searches was performed using 300,000 model points sampled from the 19-dimensional pMSSM parameter space [128]. A subsequent interpretation was performed restricting attention to a 5D sub-space of the pMSSM to specifically assess the impact of the ATLAS Run 1 searches on the electroweak production of SUSY particles, and their corresponding constraints on dark

matter [129]. The CMS collaboration also performed a reinterpretation of their Run 1 searches in the 19-dimensional pMSSM [130] using a global Bayesian analysis.

This chapter presents an interpretation of ATLAS constraints on the pMSSM using the ATLAS Run 2 searches. The study focuses specifically on the pMSSM parameters relevant to electroweakinos. Therefore, only analyses which targeted electroweakino production are considered. This ‘electroweak’ sector of the pMSSM determines much of the phenomenology relevant to dark matter, and the complementarity between collider searches and dark matter experiments will be discussed. Compared to the previous ATLAS pMSSM interpretation results in References [128] and [129], this work benefits from a broader range of search channels probed throughout the second data-taking run due to the increased statistics and centre-of-mass energy. The results provide a comprehensive overview of the status of electroweak SUSY after Run 2.

Reinterpretations of LHC SUSY searches have also been performed outside the LHC collaborations. These include results obtained using GAMBIT [131–133], MASTERCODE [134] and SMOBELS [135–137] which incorporate LHC searches into global likelihood fits in the 19-dimensional pMSSM or subsets of it. The results in this chapter benefit from accurate simulations of the ATLAS detector which are not available outside the ATLAS Collaboration. On the other hand, the computational cost of detector simulation means a smaller number of models are processed in this study compared to many truth-level studies performed outside the LHC Collaborations. This trade-off between the granularity and accuracy of the results makes reinterpretations performed within the LHC Collaborations, such as this chapter, complementary to those performed outside.

This chapter is organised as follows. Section 4.2 outlines the procedure used to sample the pMSSM and evaluate ATLAS searches on the resulting models. Two scans of the pMSSM parameter space are performed: one which randomly samples the pMSSM parameter ranges with a flat prior (named the ‘EWKino’ scan), and one which oversamples models with a viable bino-LSP dark matter candidate (named the ‘Bino-DM’ scan). The motivations and specific set-ups of each scan are outlined in Section 4.3. In Section 4.4, the sensitivity of the Run 2 ATLAS searches to the (non-simplified) pMSSM scenarios is assessed. The unique/overlapping sensitivity of the searches is also discussed. pMSSM models can also be constrained by various non-ATLAS results, including measurements of flavour and electroweak precision observables. These ‘external’ constraints are applied to the pMSSM model sets in Section 4.5, and their complementarity with ATLAS is discussed. The overall ATLAS constraints on electroweakinos for models which pass external

constraints are presented in Section 4.6. Section 4.7 discusses the complementarity between collider searches with ATLAS and direct dark matter detection experiments. The chapter concludes in Section 6.6 with a summary of the exclusion from all considered searches and constraints, a discussion of several benchmark scenarios which remain viable, and a reflection on the study performed and the future of ATLAS searches for supersymmetry.

4.2. Method

4.2.1. Overview

Before divulging the details of the methodology used in this chapter, it is instructive to provide a high-level overview. To perform a reinterpretation of ATLAS searches in the pMSSM, a sample of $\mathcal{O}(10\text{k})$ model points in the 19-dimensional pMSSM parameter space is produced via a random scan. Two scans are performed: one targeting general electroweakino production with a variety of LSP types, named the ‘EWKino scan’, and another which specifically oversamples models with a bino-like LSP that produces a viable dark matter relic density, named the ‘Bino-DM’ scan. Further details on the motivation and set-up of these scans will be provided in Section 4.3.

Proton–proton collision events are generated for each resulting model such that a suite of Run 2 searches for electroweakinos can be evaluated to determine exclusion. Table 4.1 introduces the Run 2 searches considered in this analysis, along with the simplified models targeted by each. While this does not include every Run 2 ATLAS search for electroweakinos, all key ‘wave 1’ signature searches are included. These cover every final-state lepton multiplicity that can result from various electroweakino production and decay modes. Additionally, constraints from two ATLAS Higgs analyses are included as part of the ‘ATLAS exclusion’ in this analysis. These are summarised in Table 4.2. In Section 4.4, further details of each analysis and constraint will be provided along with a discussion of their sensitivity to the EWKino and Bino-DM model sets.

In addition to the Run 2 ATLAS results, the pMSSM models can be excluded by various ‘external’ constraints. This includes flavour observables, electroweak precision measurements, the dark matter relic density, and limits from dark matter direct-detection experiments. Relevant observables are calculated for each pMSSM model such that these constraints can be applied. To allow for an evaluation of the complementarity between

Analysis	Simplified model(s) targeted
Fully Hadronic [139]	Generic heavy $\tilde{\chi}\tilde{\chi}$, including WZ , Wh and WW decays.
1Lbb [140]	Wino $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ via Wh
2L0J [141]	Wino $\tilde{\chi}_1^\pm \tilde{\chi}_1^\pm$ via WW
2L2J [142]	Wino $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ via WZ
Compressed [143]	Wino $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ via WZ , Higgsino $\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \tilde{\chi}_1^0$
3L [144]	Wino $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ via WZ , Wino $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ via Wh , Higgsino $\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \tilde{\chi}_1^0$
4L [145]	Higgsino general gauge-mediated (GGM)
Disappearing Track [146]	Wino $\tilde{\chi}_1^\pm \tilde{\chi}_1^\pm$ and $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$

Table 4.1.: ATLAS searches considered in this analysis

Analysis	Constraint
Combined Higgs boson cross-section and branching ratio measurement [147]	$m(A) > 480$ GeV
$h \rightarrow \text{inv. combination}$ [148]	$\text{BR}(h \rightarrow \text{inv.}) < 0.107$

Table 4.2.: Additional ATLAS results used to exclude pMSSM models.

ATLAS searches and these constraints, they are not applied by default to the EWKino and Bino-DM scan model sets. Instead, it will be made clear in the text which constraints have/have not been applied when producing the results in each section. There are two exceptional constraints which are applied to all models:

- Models already excluded by results from LEP [138] are removed.
- Models with a Higgs boson mass outside the window $120 \text{ GeV} < m(h) < 130 \text{ GeV}$ are removed.

All constraints considered will be discussed explicitly in Section 4.5.




Figure 4.1.: Workflow for the pMSSM scan, from sampling the pMSSM parameter space to determining if each model is excluded.

4.2.2. Scan procedure

For each scan performed, the pMSSM parameter space is randomly sampled to produce a set of model points. Each model point is then evaluated using a workflow chain summarised in Figure 4.1.

Model generation

First, the pMSSM parameter space is randomly sampled with a flat prior to produce a set of models. The specific parameter ranges used for the EWKino and Bino-DM scans will be outlined in Section 4.3. A flat prior is chosen as the relevant scan ranges are relatively narrow (< 2 TeV), with most of the interest in the range $\mathcal{O}(50 \text{ GeV})$ to $\mathcal{O}(1 \text{ TeV})$. When wider scan ranges are used, a log prior can be useful to ensure the lower end is filled. For example in Reference [149], log priors are used to scan up to 25 TeV in the context of future collider experiments.

Next, the models are passed through several programs to calculate various observables:

- SPHENO 4.0.5 β [150, 151] is used to calculate the mass spectra and decays of the SUSY particles.
- FEYNHIGGS 2.18 [152–159] is used to calculate relevant Higgs sector variables, such as the masses and branching fractions of the supersymmetric and Standard Model Higgs bosons.
- The masses of the supersymmetric particles have a significant impact on the Higgs boson masses, which FEYNHIGGS includes in its calculation. However the changes to the Higgs sector masses, in turn, influence the branching fractions of supersymmetric particles. Hence, the resulting spectrum determined with FEYNHIGGS is processed by an internal version of SPHENO provided by the authors which allows the calculation of supersymmetric branching fractions while keeping both the SUSY and Higgs sector masses constant.
- Note that the final W boson mass used as a constraint is recalculated with FEYNHIGGS 2.19.
- MICROMEAS 5.2.1 [160, 161] calculates the predicted dark matter relic density, annihilation cross-sections, and spin (in)dependent WIMP-nucleon cross-sections.

The quantities calculated by these programs are optionally used to apply constraints to the model sets which enforce consistency with current measurements of each observable. These constraints will be discussed further in Section 4.5. Models that fail one or more of these programs or contain un-physical spectra are removed. Additionally, models where the LSP is not the lightest neutralino ($\tilde{\chi}_1^0$) are removed so that every model contains a dark matter candidate.

For each model, the cross-sections for each SUSY pair-production process are calculated using PROSPINO [162]. The calculation is performed at both lead-order (LO) and next-to-leading-order (NLO). If the LO cross-section is below 2×10^{-4} pb the NLO calculation is skipped and the final cross-section is taken to be $\sigma = \sigma_{\text{LO}} \times 1.3$, with the 1.3 scale factor chosen ad-hoc to approximate the NLO contribution.

At this stage, a filter is applied to halt the processing of models that the ATLAS searches are very unlikely to have sensitivity to. The filters applied are outlined in Table 4.3. The models which do not pass these filters are included in the final model set and considered to not be excluded by the direct searches. However, their exclusion is still assessed using ATLAS constraints on $m(A)$ and $\text{BR}(h \rightarrow \text{inv.})$.

Filter	Reason
$m(\tilde{\chi}_1^\pm) < 1200$ GeV	The analyses considered have no sensitivity to scenarios with a $\tilde{\chi}_1^\pm$ mass above this limit.
$\sigma_{\text{EW}} > 7 \times 10^{-5}$ pb	If the total electroweak production cross-section is below this limit, less than 10 expected events were produced in the entire 139fb^{-1} ATLAS Run 2 dataset.
$\tilde{\chi}_1^\pm$ width > 0	The event generation procedure is not set up to handle stable charginos. These scenarios are not considered in this analysis.

Table 4.3.: Filters applied to models before event generation. Models which fail these filters are included in the final results but considered to not be excluded by the direct searches.

Truth-level evaluation

Next, events are simulated for each model which passes the filters. Producing fully reconstructed Monte Carlo (MC) samples with GEANT4 or an ATLAS fast simulation is computationally expensive, especially in this case where $\mathcal{O}(100\text{k})$ events for $\mathcal{O}(10\text{k})$ models need to be processed. Therefore initially, events are simulated at particle- (‘truth’)-level and SIMPLEANALYSIS [112] is used to obtain an approximate result that indicates if the model is likely to be excluded or not. The smearing feature of SIMPLEANALYSIS, introduced in Section 3.6.1, is applied to approximate the impact of detector effects. These results are therefore referred to as ‘truth-smearred’.

Each analysis signal region is implemented at truth-level in `SIMPLEANALYSIS` so that an estimate of signal region yields for each pMSSM model can be obtained. Validation studies of each `SIMPLEANALYSIS` implementation have been performed to confirm that the yields obtained agree with the yields for fully reconstructed samples within a reasonable margin of error (around 20%).

The particle-level simulations that are analysed using `SIMPLEANALYSIS` are performed using `MADGRAPH2.9.5` [163, 164] for process and decay generation and `PYTHIA8.306` [99] for shower generation. All electroweakino production processes are produced simultaneously in `MADGRAPH`. To reduce processing time, only a single additional jet from initial state radiation is allowed in `MADGRAPH`.

Note that some models with very small chargino decay width ($< 2 \times 10^{-14}$ GeV) fail event production. These models are removed unless they lie within the region of sensitivity of the disappearing track search [146], in which case their exclusion is determined using the disappearing track upper cross-section limits as outlined in Section 4.2.3.

The number of events generated is set to

$$N_{\text{gen.}} = \begin{cases} 10,000 & , \quad \sigma_{\text{EW}} \times \mathcal{L} \times 5 \leq 10,000 \\ 250,000 & , \quad \sigma_{\text{EW}} \times \mathcal{L} \times 5 \geq 250,000 \\ \sigma_{\text{EW}} \times \mathcal{L} \times 5 & , \quad \text{otherwise (rounded up to nearest 10,000)} \end{cases} \quad (4.1)$$

where σ_{EW} is the total electroweak SUSY production cross-section (in pb) and $\mathcal{L} = 139,000 \text{ pb}^{-1}$ (ie. the total Run 2 luminosity). The 250,000 event cap is necessary due to CPU limitations.

To improve event statistics for higher cross-section models, if $N_{\text{gen.}} > 200,000$ an additional truth-level production is processed using an event filter. The filter used is chosen to contain the event selections of all the searches and consists of a logical OR of the following:

- One lepton with $p_{\text{T}} > 15$ GeV and $|\eta| < 2.7$.
- Two leptons with $p_{\text{T}} > 3$ GeV and $|\eta| < 2.7$.
- $p_{\text{T}}^{\text{miss}} > 100$ GeV.

The number of filtered events generated is set to

$$N_{\text{gen.}}^{\text{filt.}} = \begin{cases} 10,000 & , \quad \sigma_{EW} \times \mathcal{L} \times 5 \times \epsilon_{\text{approx.}} \leq 10,000 \\ 200,000 & , \quad \sigma_{EW} \times \mathcal{L} \times 5 \times \epsilon_{\text{approx.}} \geq 200,000 \\ \sigma_{EW} \times \mathcal{L} \times 5 \times \epsilon_{\text{approx.}} & , \quad \text{otherwise (rounded up to nearest 10,000)} \end{cases} \quad (4.2)$$

where $\epsilon_{\text{approx.}}$ is set to 0.2 based on test productions that indicated that the filter efficiency is approximately 20% for most models. This is indeed the case for the majority of the models for which filtered event production is run. For a small number of models, the filter efficiency turned out to be much larger than 20%, which results in lower event statistics than expected. For some models with $N_{\text{gen.}} \approx 250,000$ and a filter efficiency of $\epsilon \gg 20\%$, the event sample produced without a filter results in smaller statistical uncertainties than the sample produced with a filter – in this case, the non-filtered sample is used for the final truth-level result. The 200,000 event cap is necessary due to CPU limitations. The number of models generated with ‘low statistics’ as a result (defined as the number of events generated is less than $2 \times \sigma_{EW} \times \mathcal{L}$) will be indicated at the end of this section in Table 4.8.

To determine if a model is excluded, a statistical fit is run using the signal yields obtained from SIMPLEANALYSIS along with an estimate of the background yields and observed data for each analysis. Most Run 2 ATLAS SUSY analyses have published a ‘likelihood’ on HEPData¹ [165]. The likelihood, introduced in Section 3.6, is a JSON file containing background and observed yield estimates for every analysis region, along with systematic uncertainty deviations. The signal yield obtained using SIMPLEANALYSIS for each analysis is “patched” into the corresponding likelihood and a hypothesis test run using PyHF [114, 115] to produce an observed and expected CL_s . A model is considered to be excluded by an analysis at truth-level if $\text{CL}_s^{\text{Obs.}} < 0.05$. Note that signal systematic uncertainties are not calculated by SIMPLEANALYSIS so are suppressed from this truth-level statistical fit.

Detector-level evaluation

The CL_s values obtained thus far are only an approximation, as truth-smearred signal samples are used and simplifications to the statistical fit are made. These approximate

¹Where this is not the case, the likelihood was obtained internally.

Category	Criteria (using truth-smeared results)
Likely excluded	$\text{CL}_s^{\text{Exp.}} < 0.001$ for at least one analysis.
Likely not excluded	$\text{CL}_s^{\text{Exp.}} > 0.1$ for every analysis.
Ambiguous	$0.001 \leq \text{CL}_s^{\text{Exp.}} \leq 0.1$ for at least one analysis and $\text{CL}_s^{\text{Exp.}} \geq 0.001$ for every analysis.

Table 4.4.: Particle-level $\text{CL}_s^{\text{Exp.}}$ criteria used to categorise models. Only models in the ‘ambiguous’ category are run through the detector-level production chain.

CL_s values are used to categorise each model as ‘likely excluded’, ‘likely not excluded’, or ‘ambiguous’. For models in the ‘likely excluded’ or ‘likely not excluded’ categories, the truth-smeared $\text{CL}_s^{\text{Obs.}}$ is used to determine exclusion. For ‘ambiguous’ models, a detector-level evaluation of each analysis is performed.

Table 4.4 presents the criteria used to categorise models as ‘likely excluded’, ‘likely not excluded’ or ‘ambiguous’ based on the truth-smeared results. These bounds are informed by comparisons of truth- and detector-level results for a test sample of models. The test sample includes models with $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ masses that lie both inside and outside the simplified model exclusion contour for each analysis, and with truth-level $\text{CL}_s^{\text{Exp.}}$ greater than and less than 0.05. This sample has been used to perform extensive validation of the procedure used, including comparisons of truth-level vs. detector-level CL_s and signal region yields for each analysis included. As a summary, Figure 4.2 compares the truth- and detector-level $\text{CL}_s^{\text{Exp.}}$ for these test models, showing the analysis with the lowest $\text{CL}_s^{\text{Exp.}}$ for each point as indicated by the colour. A bias towards smaller truth-level $\text{CL}_s^{\text{Exp.}}$ compared to detector-level is seen, which is expected due to the lack of signal systematic uncertainties and overestimation of the signal yields at truth-level. Models which lie in the red regions of Figure 4.2 would be miscategorised with the proposed decision criteria. For the test models shown, no models lie in these regions.

For models in the ‘ambiguous’ category, detector-level MC events are generated and processed with ATLFast-II simulation to obtain a more accurate result. This is computationally expensive but only needs to be done for a small proportion of the models.

The number of detector-level events generated is set to $N_{\text{reco}} = \sigma_{\text{EW}} \times \mathcal{L} \times 5 \times \epsilon$ (rounded up to the nearest 10,000) where σ_{EW} is the total electroweak SUSY production cross-section (in pb), $\mathcal{L} = 139,000 \text{ pb}^{-1}$ (ie. the total Run 2 luminosity) and ϵ is




Figure 4.2.: Comparison of detector-level and truth-level $CL_s^{\text{Exp.}}$ for a test sample of pMSSM models. The points are coloured by the analysis result being displayed, which is the analysis with the lowest $CL_s^{\text{Exp.}}$ for that model. A vertical line at 0.05 on the x -axis indicates the detector-level exclusion boundary. Horizontal lines at 0.001 and 0.1 on the y -axis indicate the final decision boundary for evaluating a model at detector-level. The red regions indicate CL_s scenarios which would result in the false categorisation of a model.

the generator filter-efficiency (known from the truth-level production²). Due to CPU limitations, any models with $N_{\text{reco}} > 300,000$ are run at truth-level only (even if they are categorised as ‘ambiguous’). In practice, this means that most low mass models (which have the highest cross-sections), particularly with $m(\tilde{\chi}_1^\pm, \tilde{\chi}_2^0) < 200$ GeV, are run at truth-level only. Additionally, models which include sparticles with decay width smaller than 2×10^{-13} or have a wino-like LSP are not processed at detector-level.

The full ATLAS analyses are then applied to the resulting MC samples using RECAST [118, 119] – a framework for re-using existing analyses to interpret new physics models that was introduced in Section 3.6.3. Finally, the CL_s output by RECAST is used to categorize ambiguous models as excluded ($\text{CL}_s^{\text{Obs.}} < 0.05$) or not excluded ($\text{CL}_s^{\text{Obs.}} > 0.05$) based on the full detector-level analysis. Multiple analyses are evaluated, and the final exclusion decision uses the observed $\text{CL}_s^{\text{Obs.}}$ for the analysis with the best (smallest) expected $\text{CL}_s^{\text{Exp.}}$. As stated previously, models with $m(A) < 480$ GeV or $\text{BR}(h \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0) > 0.107$ are also considered to be ‘excluded’ by ATLAS constraints on bSM Higgs physics.

4.2.3. Analysis approximations

To reduce CPU usage while maintaining accurate results for this large-scale workflow, some simplifications and approximations are made when evaluating each analysis. These are summarised in this section.

In cases where running the truth-level fit with the full likelihood is exceptionally slow, a simplified likelihood is used instead [116]. This makes several simplifications, such as merging all backgrounds into a single sample and reducing all nuisance parameters in the full likelihood to a single constrained parameter. This significantly increases the speed of the fit, at the price of some accuracy. A simplified likelihood is used for the Compressed, 3L (offshell) and 4L analyses, for which the full likelihood fit takes several hundred seconds.

For every analysis, signal systematic uncertainties are not calculated by SIMPLE-ANALYSIS so are suppressed from the final truth-level statistical fit. Signal statistical uncertainties are also suppressed from the final fit. As discussed previously, some model points with large cross-sections have low MC statistics which could result in false non-exclusion if signal statistical uncertainties are included. Instead, the statistical variations

²Where not known, a conservative value of $\epsilon = 0.8$ is used.

between models are expected to average out in the final binned plots. Note that all systematic and statistical uncertainties are retained in the workspace for all background samples. The detector-level evaluation includes all statistical and experimental systematic uncertainties for the signal. Note that theoretical uncertainties are not calculated for the signal so are not included in both the truth-level and detector-level evaluation.

For the Compressed analysis, poor agreement between the truth-smeared simplified likelihood fit result and the full detector-level result is found when testing simplified model points. The agreement is improved by removing low p_T^{miss} signal regions, for which SIMPLEANALYSIS does not provide an accurate approximation, from the truth-level fit. Additionally, the Compressed analysis used a MADSPIN re-weighting procedure after event generation to account for spin correlations. This procedure is not applied here. This particularly affects the shape of the di-lepton invariant mass distribution, which impacted the observed exclusion for simplified models with $M_1 \times M_2 > 0$ compared to $M_1 \times M_2 < 0$ in the original Compressed analysis – see Figures 2 and 14 of Reference [143].

The disappearing track analysis is not evaluated with SIMPLEANALYSIS or RECAST as the event generation is not set up to correctly handle long-lived particles. Instead, the upper limits on production cross-sections in the $m(\tilde{\chi}_1^\pm) - \tau(\tilde{\chi}_1^\pm)$ plane are used to determine exclusion. A linear interpolation between the available upper limits (provided internally) is used to determine the appropriate limit for each pMSSM model. To ensure a reliable interpolated result, only models with chargino mass and lifetime within the ranges for which limits are available are considered. If the total $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ production cross-section for a model, $\sigma(\tilde{\chi}_1^\pm \tilde{\chi}_2^0) + \sigma(\tilde{\chi}_1^\pm \tilde{\chi}_1^\pm)$, is greater than the corresponding interpolated limit then the model is considered excluded. Only direct $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ production is considered which leads to slightly conservative limits as production through decay of heavier states is ignored. Note that the limits used were produced in a grid of $\tilde{\chi}_1^\pm$ mass and lifetime with constant values of $\tan \beta$ and μ , but are applied here to pMSSM scenarios with varying $\tan \beta$ and μ . Therefore, the evaluation procedure uses the approximation that the disappearing track analysis acceptance is determined by the $\tilde{\chi}_1^\pm$ mass and lifetime only and is independent of model parameters such as $\tan \beta$ or μ .

For some analyses, the truth-smeared yields are found to consistently underestimate detector-level yields by a scale factor. This is likely due to outdated or incorrect efficiency factors in the SIMPLEANALYSIS smearing. In this case, a scaling factor is calculated via a least-squares fit of the truth-smeared and detector-level yields for a test sample of pMSSM points. This scaling factor is then applied to all truth-smeared yields for that analysis.

Analysis	Simplifications and approximations made
FullHad	Truth-smeared yields are scaled by 0.62
1Lbb	Truth-smeared yields are scaled by 0.81
2L0J	Truth-smeared yields are scaled by 0.80
2L2J	• Truth-smeared yields are scaled by 0.70 • Signal contamination in control regions neglected at truth-level.
Compressed	• Simplified likelihood used at truth-level. • Low p_T^{miss} signal regions removed at truth-level. • No MADSPIN re-weighting applied. • ISR uncertainties are not included.
3L (onshell)	–
3L (offshell)	Simplified likelihood used at truth-level.
4L	Simplified likelihood used at truth-level.
Disappearing track	The analysis is not evaluated with SIMPLEANALYSIS or RECAST. Instead, the upper cross-section limits are applied.

Table 4.5.: Approximations and simplifications made when evaluating each analysis

The truth-level evaluation of the 2L2J analysis neglects signal contamination in the control regions. All other truth-level statistical fits, and all detector-level statistical fits, do include an estimate of the signal yields in the control regions to account for signal contamination.

A summary of specific simplifications made for each analysis is shown in Table 4.5. The 3L (onshell) analysis is the only analysis for which no approximations are required.

As discussed above, models with a large cross-section are processed at truth-level only due to CPU limitations. Table 4.6 shows the impact of this restriction by comparing the total number of models that lie within the detector-level selection criteria for each analysis to the number of models that are actually processed at detector-level. The Compressed and 3L (offshell) analyses, which cover the low-mass compressed region, are the most severely impacted. Models with a wino-like LSP are also not considered for

Analysis	Number of models in detector-level selection			
	EWKino scan		Bino scan	
	All	Processed	All	Processed
FullHad	454	391	52	46
1Lbb	178	157	231	220
2L0J	62	47	10	5
2L2J	269	197	630	447
Compressed	193	7	161	8
3L (onshell)	227	141	755	455
3L (offshell)	80	13	591	182
4L	12	6	16	10
Any	987	770	1440	859

Table 4.6.: Number of models that pass the detector-level selection for each analysis (from Table 4.4), and the number that are actually processed at detector-level. Models with a wino-like LSP were not considered for detector-level production so are omitted from these counts.

detector-level evaluation. The disappearing track search is the most sensitive search to wino-LSP scenarios but does not require event production as it is evaluated using upper limits on production cross-sections as discussed above.

4.3. A tale of two scans

The workflow described in the previous section is used to process two pMSSM scans with different parameter ranges and sampling strategies. The details and motivation for each scan will be outlined here.

4.3.1. EWKino scan

The first scan focuses on general electroweakino production, with other sparticles decoupled. This scan will be referred to as the ‘EWKino’ scan. The EWKino scan includes 15,537 models which are successfully generated with a $\tilde{\chi}_1^0$ LSP from an initial set of 20,000 models randomly sampled (with a flat prior) from the ranges listed in Table 4.7.

Parameter	min	max	Note
$m_{\tilde{L}_1} (=m_{\tilde{L}_2})$	10 TeV	10 TeV	Left-handed slepton (first two gens.) mass
$m_{\tilde{e}_1} (=m_{\tilde{e}_2})$	10 TeV	10 TeV	Right-handed slepton (first two gens.) mass
$m_{\tilde{L}_3}$	10 TeV	10 TeV	Left-handed stau doublet mass
$m_{\tilde{e}_3}$	10 TeV	10 TeV	Right-handed stau mass
$m_{\tilde{Q}_1} (=m_{\tilde{Q}_2})$	10 TeV	10 TeV	Left-handed squark (first two gens.) mass
$m_{\tilde{u}_1} (=m_{\tilde{u}_2})$	10 TeV	10 TeV	Right-handed up-type squark (first two gens.) mass
$m_{\tilde{d}_1} (=m_{\tilde{d}_2})$	10 TeV	10 TeV	Right-handed down-type squark (first two gens.) mass
$m_{\tilde{Q}_3}$	2 TeV	5 TeV	Left-handed squark (third gen.) mass
$m_{\tilde{u}_3}$	2 TeV	5 TeV	Right-handed top squark mass
$m_{\tilde{d}_3}$	2 TeV	5 TeV	Right-handed bottom squark mass
M_1	-2 TeV	2 TeV	Bino mass parameter
M_2	-2 TeV	2 TeV	Wino mass parameter
μ	-2 TeV	2 TeV	Bilinear Higgs mass parameter
M_3	1 TeV	5 TeV	Gluino mass parameter
A_t	-8 TeV	8 TeV	Trilinear top coupling
A_b	-2 TeV	2 TeV	Trilinear bottom coupling
A_τ	-2 TeV	2 TeV	Trilinear τ lepton coupling
M_A	0 TeV	5 TeV	Pseudoscalar Higgs boson mass
$\tan \beta$	1	60	Ratio of the Higgs vacuum expectation values

Table 4.7.: Parameter ranges which are randomly sampled for the ‘EWKino’ scan.

The slepton, stau and first-generation squark mass parameters are set to 10 TeV such that they are decoupled. The third-generation squark mass parameters and gluino mass parameter are kept lighter to increase the number of models with a Standard Model-like Higgs boson mass. They are allowed to vary between 2–5 TeV and 1–5 TeV respectively. Note that gluino and squark production is not generated, and the resulting model exclusion from the electroweak analyses is found to be independent of these parameters.

The electroweakino masses are dominantly controlled by the bino, wino and bilinear Higgs mass parameters. ATLAS sensitivity to electroweakinos only extends to 1 TeV, and

these parameters are varied between -2 TeV and $+2$ TeV to ensure a large proportion of models are within the region of ATLAS sensitivity while still allowing for scenarios with large mass splittings between the lighter and heavier electroweakinos. Reducing or increasing this range would impact the relative fraction of models that the ATLAS searches are sensitive to, which one should keep in mind when interpreting the results.

The trilinear top-quark coupling parameter A_t has a significant impact on the loop corrections to the Higgs boson mass, with a larger $|A_t|$, in general, leading to a Higgs boson mass that is closer to the experimentally measured value for a given model point. Therefore a large range of -8 TeV to $+8$ TeV is used for this parameter. A smaller range of -2 TeV to $+2$ TeV is used for the bottom-quark and τ -lepton trilinear couplings.

The pseudoscalar Higgs boson mass impacts the decay modes of the electroweakinos (which can decay to A if it is lighter). It also affects the dark matter phenomenology for each model point – a lighter A mass allows for an additional co-annihilation mechanism $\tilde{\chi}_1^0 \tilde{\chi}_1^0 \rightarrow A$ which can reduce the dark matter relic density. To cover a range of phenomenology from light to decoupled A , a wide range of 0 to 5 TeV is chosen. Finally, $\tan \beta$ (the ratio of the vacuum expectation values of the two Higgs doublets) is varied between 1 and 60.

Figure 4.3a shows the models from the EWKino scan in the plane of $m(\tilde{\chi}_1^0)$ vs. Ωh^2 (the dark matter relic density), coloured with RGB value set to the wino/bino/higgsino fraction of the LSP. The observed dark matter relic density measured by the Planck mission [33], $\Omega h^2 = 0.12$, is indicated by a horizontal dashed line. A higgsino- or wino-like LSP underestimates the relic density, unless the LSP mass is $\mathcal{O}(1$ TeV). This is because a higgsino- or wino-like LSP has very efficient co-annihilation with the nearby $\tilde{\chi}_1^\pm$ and/or $\tilde{\chi}_2^0$. This is demonstrated in Figure 4.3b which colours each model by the dominant annihilation mechanism. For the higgsino- and wino-like LSP models, $\tilde{\chi}_1^\pm / \tilde{\chi}_2^0$ co-annihilation is the dominant mechanism. On the other hand, a bino-like LSP, in general, overestimates the relic density – in Figure 4.3a, almost all bino-like LSP models lie above the dashed line. Note that if the squark or slepton masses were lower than the values in Table 4.7 then squark or slepton co-annihilation could also alter the relic density.

The LSP does not necessarily saturate the observed relic density. Dark matter could have multiple components with an additional hidden sector contributing to the relic density. Therefore the Planck measurement is considered an upper limit, such that models which predict $\Omega h^2 \leq 0.12$ are considered viable. This includes almost all models




Figure 4.3.: Scatter of models from the EWKino scan in the $m(\tilde{\chi}_1^0)$ vs. Ωh^2 plane. (a) is coloured with RGB value set to the wino/higgsino/bino fraction of the LSP and (b) is coloured by the dominant annihilation mechanism. The experimentally measured value of the relic density [33], $\Omega h^2 = 0.12$, is indicated by a horizontal dashed line.

with a higgsino- or wino-like LSP, but very few models with a bino-like LSP. Additionally, alternatives to the thermal freeze-out mechanism have been proposed which may allow models which have a calculated relic density greater than 0.12 here [166]. Therefore to avoid being overly restrictive, for analysis of the EWKino scan, the dark matter relic density is not used to constrain the models. Instead, a separate scan, named the ‘Bino-DM’ scan, is performed to oversample models with a bino-like LSP which do satisfy $\Omega h^2 \leq 0.12$ and allow an evaluation of a bino-like LSP as a dark matter candidate within the thermal freeze-out paradigm.

4.3.2. Bino-DM scan

The Bino-DM scan provides a sample of bino-LSP models which satisfy $\Omega h^2 \leq 0.12$. The ranges of this scan are the same as those in Table 4.7 except the bino mass parameter range is tightened to $|M_1| < 500$ GeV to focus on scenarios with a bino-like LSP. A random sample of 437,500 models from these ranges is produced, and only those with a




Figure 4.4.: Scatter plot of models from the Bino-DM scan in the $m(\tilde{\chi}_1^0)$ vs. Ωh^2 plane. (a) is coloured by the dominant annihilation mechanism and (b) is coloured by the LSP higgsino fraction with black representing 0. The experimentally measured value of the relic density [33], $\Omega h^2 = 0.12$, is indicated by a horizontal dashed line.

bino-like LSP (defined as bino fraction greater than 0.5) and $\Omega h^2 \leq 0.12$ are kept. These constraints reduce the number of models to 11,163.

Figure 4.4a shows the initial sample of bino-LSP models prior to requiring $\Omega h^2 \leq 0.12$, coloured by the dominant dark matter annihilation mechanism. Several regions of interest can be observed. First, the so-called ‘Z/h-funnel’ regions are coloured purple. When the LSP mass is approximately half the mass of the Z or Higgs boson, the self-annihilation into a Z or Higgs boson is enhanced which reduces the dark matter relic density to an acceptable value. To couple to the h and Z bosons such that these annihilation mechanisms are allowed, the LSP must have a higgsino component. This can be seen in Figure 4.4b where these bino-LSP models are coloured by the higgsino fraction of the LSP. Around the funnel regions, a significant LSP higgsino fraction is seen. This is particularly noticeable on the edges of the funnels. The lower the higgsino fraction, the smaller the coupling between the LSP and h/Z, so the closer $m(\tilde{\chi}_1^0)$ has to be to the $m(Z)/2$ or $m(h)/2$ resonance to enhance the self-annihilation sufficiently to bring the relic density below 0.12.

The funnel regions are of particular interest as they can produce a satisfactory relic density with an LSP mass which overlaps with the region of sensitivity of many of the Run 2 ATLAS searches for SUSY. Similarly, there are A and H funnel regions where $m(\tilde{\chi}_1^0) \approx m(A, H)/2$, corresponding to enhanced LSP self-annihilation into the second scalar Higgs boson H or pseudoscalar Higgs boson A .

For $m(\tilde{\chi}_1^0) > 100$ GeV, in the so-called ‘bulk region’, a variety of other annihilation mechanisms contribute. When $m(\tilde{\chi}_1^0) > 173$ GeV (ie. the LSP is greater in mass than the top quark), the $\tilde{\chi}_1^0 \tilde{\chi}_1^0 \rightarrow t\bar{t}$ self-annihilation process is allowed, coloured in dark pink. The orange points also show models which have a nearly pure bino LSP which is close in mass to a nearly pure wino $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$. In these scenarios, $\tilde{\chi}_1^\pm$ - $\tilde{\chi}_2^0$ annihilation competes with their decays to the LSP. The relic density is reduced as the $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ annihilate into SM particles before they can decay into the LSP.

In summary, requiring $\Omega h^2 \leq 0.12$ and a bino-like LSP favours two spectra which will be particularly relevant to the interpretations in Section 4.6. The first is a bino-like LSP with a small higgsino component in the funnel regions. This requires that the higgsino mass is relatively light to allow the mixing, and therefore the $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ are also likely to be higgsino-like. This is unlike typical simplified models of electroweak production which assume a pure-bino LSP and pure-wino $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$. The electroweakino mass-spectrum for such a model is shown in Figure 4.5a. The second is a pure-bino LSP which is close in mass to a wino-like $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$. This scenario is closer to the typical simplified models considered by direct searches in ATLAS. The electroweakino mass-spectrum for such a model is shown in Figure 4.5b.

4.3.3. Summary of the scans

Table 4.8 summarises the names, bino mass parameter ranges, and the number of models processed at each workflow step, for the two scans. For both scans, models which are already excluded by results from LEP [138] or have a Higgs boson mass outside the window $120 \text{ GeV} < m(h) < 130 \text{ GeV}$ are removed before producing the final results shown in the following sections. The number of models passing these requirements is also shown in Table 4.8.

Scan name	EWKino	Bino-DM
$ M_1 $ range	0 – 2 TeV	0 – 500 GeV
LSP type	Neutralino	Bino-like neutralino
Number of models generated:		
Sampled	20,000	437,500
Successful generation	16,883	370,058
Correct LSP type	15,537	286,308
Pass DM constraint	N/A	11,163
Number of models processed:		
Pass pre-filter	12,704	11,007
Successful truth-level evaluation	12,440	10,948
Low statistics truth-production	1,604	3,589
Reco-level evaluation	770	859
Number of models:		
Pass LEP constraint	13,931	10,165
120 GeV < $m(h)$ < 130 GeV	12,280	8,897

Table 4.8.: Set-ups for the two pMSSM scans performed and the number of models passing each step. ‘Low statistics truth-production’ here is defined as the number of events generated is less than $2 \times \sigma_{\text{EW}} \times \mathcal{L}$.




Figure 4.5.: Mass-spectrum for two Bino-DM scan models: (a) A model in the funnel region with a bino/higgsino mixed LSP and higgsino-like $\tilde{\chi}_1^\pm$, (b) A model in the bulk region with an almost pure-bino LSP and pure-wino $\tilde{\chi}_1^\pm$. Produced using PySLHA [167].

4.4. The Run 2 ATLAS searches for electroweak supersymmetry

In this section, the Run 2 ATLAS searches for electroweak supersymmetry used in this analysis are presented. These searches were based on 139 fb^{-1} of proton–proton collisions at $\sqrt{s} = 13$ TeV recorded by the ATLAS detector during Run 2 of the LHC. No statistically significant signals were observed in these searches and, instead, each reported exclusion limits on electroweakinos in particular simplified models.

Each search targets either $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_1^\pm$ or $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_2^0$ pair-production, followed by a specific decay mode. The $\tilde{\chi}_1^\pm$ decays to the LSP via $\tilde{\chi}_1^\pm \rightarrow W^\pm \tilde{\chi}_1^0$, where the W -boson can subsequently decay leptonically ($W^\pm \rightarrow \ell^\pm \nu_\ell$) or hadronically ($W^\pm \rightarrow q\bar{q}'$). The second neutralino on the other hand can decay to the LSP via $\tilde{\chi}_2^0 \rightarrow h\tilde{\chi}_1^0$, $\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0$ and $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^\pm W^\pm$. Subsequent h , Z and W decays also produce leptons or hadronic jets. Depending on the specific production mechanism and subsequent decay mode targeted, each analysis' signal selection includes 0–4 leptons and 0–4 hadronic jets.

For each search, the targeted simplified model(s) are outlined along with relevant key features of the analysis strategy. Then, the sensitivity of the search to pMSSM scenarios will be presented and compared to the original exclusion reported for simplified models. A model is considered ‘excluded’ by an analysis if the final observed $\text{CL}_s^{\text{Obs.}}$ is

less than 0.05. In general, the EWKino scan will be used in this section, with the goal of exploring the sensitivity of each analysis to non-simplified scenarios. The EWKino scan fills the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ and $\tilde{\chi}_2^0 - \tilde{\chi}_1^0$ mass planes in which most of the original analyses reported their simplified model exclusion, so is useful for comparing the pMSSM sensitivity to the original published contours. Where the sensitivity to the Bino-DM scan is significantly different, this will be discussed as well.

4.4.1. Direct searches

Fully hadronic

The fully hadronic analysis [139], referred to as ‘FullHad’, searched for pair-production of heavy electroweakinos, with decays via W , Z , or Higgs bosons to the LSP. Unlike many of the other analyses included, FullHad targeted a generic signature rather than assuming particular electroweakinos are purely wino/higgsino/bino. Diagrams of the generic production processes considered are shown in Figure 4.6. Here $\tilde{\chi}_{\text{heavy}}$ can be a wino- or higgsino-like chargino or heavy neutralino, and $\tilde{\chi}_{\text{light}}$ is the LSP (in our case $\tilde{\chi}_1^0$) and can be bino-, wino- or higgsino-like.

Figure 4.7a shows the fraction of pMSSM models excluded by FullHad projected onto the two-dimensional plane of the masses of the lightest chargino and LSP. For comparison, the observed 95% CL exclusion limit on a simplified model of pure-wino $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W \tilde{\chi}_1^0 Z \tilde{\chi}_1^0$ and $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_1^\pm \rightarrow W \tilde{\chi}_1^0 W \tilde{\chi}_1^0$ with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z \tilde{\chi}_1^0) = 100\%$




Figure 4.6.: Diagrams of the signatures targeted by the fully hadronic (FullHad) analysis.






Figure 4.7.: (a) Fraction of models excluded by FullHad in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane for the EWKino scan. Grey bins contain no models and black bins have 0% exclusion. The observed exclusion contour on a simplified model of pure-wino $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W\tilde{\chi}_1^0 Z\tilde{\chi}_1^0$ and $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_1^\pm \rightarrow Z\tilde{\chi}_1^0 Z\tilde{\chi}_1^0$ with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0) = 100\%$ and $\text{BR}(\tilde{\chi}_1^\pm \rightarrow W\tilde{\chi}_1^0) = 100\%$ is overlaid. (b) The distribution of the $\tilde{\chi}_1^\pm$ wino fraction for models which lie inside the FullHad simplified model exclusion contour, shown before and after applying FullHad. The ratio panel shows the fraction of models excluded in each bin.

and $\text{BR}(\tilde{\chi}_1^\pm \rightarrow W\tilde{\chi}_1^0) = 100\%$ is overlaid. The pattern of sensitivity matches the area covered by the contour well.

Figure 4.7b shows the 1D distribution of the $\tilde{\chi}_1^\pm$ wino fraction before and after applying the FullHad exclusion. To isolate models for which FullHad is expected to have sensitivity, only models which lie inside the FullHad simplified model exclusion contour are shown. There is a strong preference for wino-like $\tilde{\chi}_1^\pm$ (as opposed to higgsino-like) as this corresponds to a larger production cross-section – this is shown in Figure 4.8. No clear preference for a particular $\tilde{\chi}_2^0$ decay mode is found: Of the pMSSM models which lie inside the FullHad simplified model exclusion contour, 25% of the models with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0) > 50\%$ are excluded and 29% of the models with $\text{BR}(\tilde{\chi}_2^0 \rightarrow h\tilde{\chi}_1^0) > 50\%$ are excluded.

FullHad also has sensitivity to some models on the diagonal in Figure 4.7a, where the $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_1^0$ are close in mass. Figure 4.9a shows this more clearly, with $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ on the y -axis. This is due to FullHad’s sensitivity to generic production of heavy electroweakinos, which in the pMSSM includes the heavier charginos and neutralinos, $\tilde{\chi}_{2,3}^\pm$ and $\tilde{\chi}_{3,4}^0$, which



can decay to the lighter $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ via Z and W bosons. This is demonstrated in Figure 4.9b which shows the total truth-smeared signal region yields for models which are excluded at truth-level by FullHad and satisfy $m(\tilde{\chi}_1^\pm) < 400$ GeV, $m(\tilde{\chi}_1^0) < 400$ GeV and $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) < 50$ GeV, binned by SUSY production process. $\tilde{\chi}_2^\pm \tilde{\chi}_2^\pm$ and $\tilde{\chi}_3^0 \tilde{\chi}_2^\pm$ are the dominant production mechanisms. In general, the production of heavier charginos and neutralinos will lead to too many jets from the subsequent multi-step decay chain, so the FullHad signal selections will not be satisfied. However, when $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_{2,1}^0$ are compressed, they effectively form a collective p_T^{miss} system as the $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ decay products are so soft as to be unobservable. Hence, only the decay products of the heavier electroweakinos are observable, and the correct number of jets to pass the selections can be produced.

1Lbb

The ‘1Lbb’ analysis [168] targeted $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_2^0$ pair-production where the second neutralino, $\tilde{\chi}_2^0$, decays into a Standard Model Higgs boson and the LSP ($\tilde{\chi}_1^0$). The diagram of this process is shown in Figure 4.10. The following simplified model assumptions were made:

- The LSP is pure bino.
- $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ are pure wino.
- $\text{BR}(\tilde{\chi}_2^0 \rightarrow h\tilde{\chi}_1^0) = 100\%$
- $\text{BR}(\tilde{\chi}_1^\pm \rightarrow W\tilde{\chi}_1^0) = 100\%$
- $\text{BR}(h \rightarrow b\bar{b}) = 58.3\%$ (expected value for the SM Higgs boson)

The 1Lbb signal selection requires one lepton from the W boson decay, plus two b -jets from the Higgs boson decay (with invariant mass consistent with the SM Higgs boson mass). The signal regions are binned in the transverse mass m_T and contranverse mass m_{CT} [169, 170] which separate the signal from W +jets and $t\bar{t}$ backgrounds, respectively.

Figure 4.11a shows the fraction of EWKino scan models excluded by 1Lbb in the $\tilde{\chi}_2^0$ – $\tilde{\chi}_1^0$ mass plane, overlain with the 95% CL observed exclusion limits on the simplified model in Figure 4.10. The area of sensitivity to the pMSSM models agrees well with the simplified model contour, though with $< 50\%$ of models excluded in most bins. This is partly explained by Figure 4.11b, which shows the distribution of $\text{BR}(\tilde{\chi}_2^0 \rightarrow h\tilde{\chi}_1^0)$ for




Figure 4.10.: Diagram of the signature targeted by the 1Lbb analysis.




Figure 4.11.: (a) Fraction of EWKino scan models excluded by 1Lbb in the $\tilde{\chi}_2^0$ - $\tilde{\chi}_1^0$ mass plane, overlain with the observed exclusion contour for the simplified model in Figure 4.10. Grey bins contain no models and black bins have 0% exclusion. (b) Distribution of $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h)$ before and after 1Lbb for EWKino models which lie inside the observed 1Lbb exclusion contour in the $\tilde{\chi}_2^0$ - $\tilde{\chi}_1^0$ mass plane. The ratio plot shows the fraction of models excluded by 1Lbb in each bin.

Model 4.12a	Model 4.12b	Model 4.12c
$\tilde{\chi}_1^0$ is pure bino		
$\tilde{\chi}_1^\pm$ is pure wino		
$\text{BR}(\tilde{\chi}_1^\pm \rightarrow W \tilde{\chi}_1^0) = 100\%$	$\text{BR}(\tilde{\chi}_1^\pm \rightarrow \tilde{\ell} \nu) = 50\%$ $\text{BR}(\tilde{\chi}_1^\pm \rightarrow \tilde{\nu} \ell) = 50\%$	$\text{BR}(\tilde{\ell} \rightarrow \ell \tilde{\chi}_1^0) = 100\%$
	$m(\tilde{\ell}) = m(\tilde{\nu}) = (m(\tilde{\chi}_1^\pm) + m(\tilde{\chi}_1^0))/2$	

Table 4.9.: Simplified model assumptions corresponding to the models for the 2L0J analysis in Figure 4.12.

models which lie within the simplified model exclusion contour. As expected, a clear correlation between this branching ratio and the 1Lbb exclusion can be seen. In the pMSSM, the $h \rightarrow b\bar{b}$ branching ratio can also deviate from the Standard Model value of 58.4% which impacts the 1Lbb acceptance. Additionally, 1Lbb exclusion is only seen for models with a wino-like $\tilde{\chi}_2^0$ due to the larger production cross-sections compared to higgsino-like $\tilde{\chi}_2^0$.

2L0J

The ‘2L0J’ analysis [60] searched for the production of charginos and sleptons decaying into final states with two leptons and missing transverse momentum. The search was optimized and interpreted using three simplified models, with diagrams shown in Figure 4.12. The assumptions made for each simplified model are summarised in Table 4.9. Since sleptons are decoupled in the EWKino and Bino-DM scans, the relevant simplified model is Figure 4.12a – pair-production of the lightest chargino. Note that in the EWKino and Bino-DM scans the third simplified model assumption, $\text{BR}(\tilde{\chi}_1^\pm \rightarrow W \tilde{\chi}_1^0) = 100\%$, is satisfied for all models³.

In addition to requiring two leptons, zero jets, and high p_T^{miss} , the key kinematic variable in this analysis was the stransverse mass: m_{T2} [95, 96]. The 2L0J signal regions are binned in m_{T2} , which has an endpoint related to the masses of the slepton/chargino and neutralino in the decays in Figure 4.12. This separates the signal from SM background processes for which the m_{T2} distribution tails off above the W boson mass.

³Though the $\tilde{\chi}_1^\pm \rightarrow W \tilde{\chi}_1^0$ decay will proceed via an offshell W boson in cases where $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) < m_W$.




Figure 4.12.: Diagrams of the signatures targeted by the 2L0J analysis.

Figure 4.13 shows the fraction of EWKino scan models excluded by 2L0J in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane. The 95% CL observed exclusion limit for the simplified model in Figure 4.12a is overlain. Good sensitivity to the pMSSM models is observed, with greater than 50% exclusion in many bins, and some exclusion beyond the reported simplified model limits. The EWKino models in this region mostly contain either a pure-wino $\tilde{\chi}_1^\pm$ or pure higgsino $\tilde{\chi}_1^\pm$ (rather than a mixture). Only models with a wino $\tilde{\chi}_1^\pm$ are found to be excluded by 2L0J, due to the larger corresponding production cross-section as previously discussed.

In the Bino-DM scan, the models in the 2L0J region of sensitivity are funnel-region scenarios which have a higgsino-like $\tilde{\chi}_1^\pm$, as shown in Figure 4.14a and discussed in Section 4.3. Consequently, the sensitivity of 2L0J to models in the Bino-DM scan, shown in Figure 4.14b, is much weaker than the EWKino scan. A similar argument applies to other analyses which cover the funnel regions of the Bino-DM scan.




Figure 4.13.: Fraction of models from the EWKino scan excluded by 2L0J in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane. Grey bins contain no models and black bins have 0% exclusion. The observed exclusion contour for the simplified model in Figure 4.12a is overlain.




Figure 4.14.: Two plots in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane: (a) shows each Bino-DM scan model coloured by the $\tilde{\chi}_1^\pm$ higgsino fraction, (b) shows the fraction of Bino-DM models excluded by 2L0J with the 2L0J simplified model exclusion contour overlain. Grey bins contain no models and black bins have 0% exclusion.

Higgsino	Wino/Bino
$(\tilde{\chi}_1^0, \tilde{\chi}_1^\pm, \tilde{\chi}_2^0)$ are a higgsino triplet	$\tilde{\chi}_1^0$ is pure bino, $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ are pure wino
$\text{BR}(\tilde{\chi}_1^\pm \rightarrow W^* \tilde{\chi}_1^0) = 100\%$ $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^* \tilde{\chi}_1^0) = 100\%$	$\text{BR}(\tilde{\chi}_1^\pm \rightarrow W^* \tilde{\chi}_1^0) = 100\%$ $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^* \tilde{\chi}_1^0) = 100\%$
$m(\tilde{\chi}_1^\pm) = (m(\tilde{\chi}_2^0) + m(\tilde{\chi}_1^0))/2$	$m(\tilde{\chi}_1^\pm) = m(\tilde{\chi}_2^0)$

Table 4.10.: Simplified model assumptions corresponding to the models for the Compressed analysis in Figure 4.15.





Figure 4.15.: Diagrams of the signatures targeted by the compressed analysis.

Compressed

The ‘compressed’ analysis [171] targets scenarios with a compressed mass-spectrum. The relevant simplified models⁴ targeted by the Compressed analysis are shown in Table 4.10. Both pair-production (Figure 4.15a) and VBF (Figure 4.15b) production modes are considered.

In both scenarios, two leptons are produced from the decay of the second neutralino to the LSP via an offshell Z boson. A very high p_T jet is also required in the form of either initial state radiation (ISR) or forward jets (for VBF production). This boosts the SUSY system, allowing higher lepton p_T thresholds and large p_T^{miss} to discriminate signal from background. In some signal regions, one lepton is assumed to be too low p_T to be reconstructed and, instead, a requirement of one lepton plus one track is imposed.

⁴Slepton models were also targeted but are not considered here.



(a) EWKino scan (higgsino-like LSP)

Figure 4.16.: Fraction of EWKino scan models with a higgsino-like LSP that are excluded by the Compressed analysis, shown in the $m(\tilde{\chi}_2^0)$ vs. $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane. Grey bins contain no models and black bins have 0% exclusion. The observed higgsino simplified model exclusion contour is overlain.


(a) $M_1 \times M_2 > 0$ (b) $M_1 \times M_2 < 0$

Figure 4.17.: Fraction of Bino-DM scan models that are excluded by the Compressed analysis, shown in the $m(\tilde{\chi}_2^0)$ vs. $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane. (a) shows models with $M_1 \times M_2 > 0$ and (b) shows models with $M_1 \times M_2 < 0$. Grey bins contain no models and black bins have 0% exclusion. The plots are overlain with the relevant observed simplified model exclusion contour for each case.

The Compressed analysis exclusion of EWKino scan models with a higgsino-like LSP is shown in Figure 4.16 overlain with the observed exclusion contour for the higgsino simplified model in Table 4.10. The exclusion is relatively poor, with less than 50% exclusion in most bins. Some exclusion outside the reported contour is observed. Note that these low-mass models are the most susceptible to statistical fluctuation, which may explain this leakage.

Stronger sensitivity is seen for models with a bino-like LSP and wino-like $\tilde{\chi}_1^\pm / \tilde{\chi}_2^0$. The EWKino scan does not contain a large sample of such models, so the Bino-DM scan is shown in Figure 4.17. Due to observed excesses/deficits in different bins of the di-muon invariant mass, the Compressed analysis saw distinct observed exclusion patterns for bino-LSP models with $M_1 \times M_2 > 0$ and $M_1 \times M_2 < 0$. Hence, each of these cases are shown separately. The exclusion follows a similar pattern to the overlain contours despite the fact that, as discussed in Section 4.2.3, the MADSPIN re-weighting procedure is not applied to the pMSSM samples.

For both higgsino- and bino-LSP scenarios, the sensitivity decreases rapidly towards the edges of the overlain contours. The difference in exclusion compared to the simplified models can partly be understood by considering the branching fraction for the $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \ell^+ \ell^-$ decay mode – which the Compressed analysis selections target.

Figure 4.18a shows a scatter of the mass difference between $\tilde{\chi}_2^0$ and the LSP versus $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \mu^+ \mu^-)$ for all models in the EWKino and Bino-DM scans that lie inside the relevant Compressed analysis exclusion contour for higgsino or wino/bino simplified models. Points are coloured with the RGB value set to the $\tilde{\chi}_2^0$ wino/higgsino/bino fraction. The green crosses show the higgsino simplified model branching ratios calculated by SUSYHIT [172] for use in the Compressed analysis. These agree well with the SPHENO values for the higgsino-LSP pMSSM models (shown as green points).

However, the simplified models considered by the Compressed analysis did not include the radiative decay $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 + \gamma)$. This decay mode was removed and the $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \ell^+ \ell^-$ branching ratios re-scaled. The radiative decay branching ratio, shown in Figure 4.18b, is significant for the pMSSM models, particularly at small mass splittings. The grey crosses in Figure 4.18a show the re-scaled $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \mu^+ \mu^-$ branching ratios which were used for the Compressed analysis simplified models. These are much larger than the branching ratios for the higgsino pMSSM models.




Figure 4.18.: Scatter plot of the mass difference between the second and lightest neutralino versus (a) $BR(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \mu^+ \mu^-)$, (b) $BR(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 + \gamma)$. The plots include pMSSM models from both the EWKino and Bino-DM scans that lie inside the relevant higgsino or wino/bino simplified model expected exclusion contour from the Compressed analysis. Points are coloured with RGB value set to the $\tilde{\chi}_2^0$ wino/higgsino/bino fraction. The crosses in (a) indicate the higgsino-LSP simplified model branching ratio from SUSYHIT (green), the branching ratio used in the Compressed analysis for higgsino-LSP simplified models (grey), and the branching ratio used in the Compressed analysis for wino-LSP simplified models (red).




Figure 4.19.: Diagram of the signature targeted by the 2L2J analysis.

Additionally, the simplified wino/bino models considered in the Compressed analysis used $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \mu^+ \mu^-) = 3.5\%$ for all points (indicated by the red crosses in Figure 4.18a). This is larger than the wino pMSSM model branching ratios (shown in red).

Overall, the more realistic pMSSM models have a smaller $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \mu^+ \mu^-$ branching ratio than the simplified models used in the original Compressed analysis interpretation. The same conclusions hold for the $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 e^+ e^-$ decay mode. This explains the reduced sensitivity to the pMSSM models compared to the simplified model exclusion contours seen in Figures 4.16 and 4.17.

2L2J

The ‘2L2J’ analysis [142] searched for SUSY production processes with a final state including two leptons, two jets and missing transverse momentum. Several simplified models were used to optimize the search strategy and interpret the results, including strong-SUSY and gauge-mediated SUSY breaking (GMSB) models. The model relevant to the pMSSM scans considered here is shown in Figure 4.19 and consists of pair-production of $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$. The following simplified model assumptions were made:

- $\tilde{\chi}_1^0$ is pure bino
- $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ are pure wino and have equal mass
- $\text{BR}(\tilde{\chi}_1^\pm \rightarrow W^{(*)} \tilde{\chi}_1^0) = 100\%$
- $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^{(*)} \tilde{\chi}_1^0) = 100\%$

For cases where $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0) > m_Z$, the two leptons originate from an onshell Z boson, so are required to be same-flavour opposite-sign and have a dilepton invariant mass consistent with the Z boson mass. Furthermore, the dijet invariant mass is required to be consistent with either the W or Z boson mass. Additional ‘offshell’ signal regions target scenarios with smaller mass splittings where the $\tilde{\chi}_2^0$ decay proceeds via an offshell Z boson. These regions instead use high p_T jets and tighter m_{T2} selections to discriminate signal from background.

The 2L2J exclusion of EWKino scan models in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane is shown in Figure 4.20a. The area of exclusion agrees well with the overlain 2L2J simplified model exclusion contour. There is leakage beyond the contour in the compressed region, which can also be seen in Figures 4.20b and 4.20c where the exclusion is shown with $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ and $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ on the y -axis, respectively.

The exclusion of compressed-mass scenarios is mainly due to contributions from heavier electroweakinos. Figure 4.21a demonstrates this, showing the total truth-smearred signal region yields for each SUSY production process, summed over the models in the EWKino scan which are excluded by 2L2J (at truth-level) and lie outside the simplified model expected exclusion contour with $m(\tilde{\chi}_1^\pm) < 300$ GeV. The yields for the ‘High’, ‘Intermediate’ and ‘Low’ p_T^{miss} signal regions, and the ‘Offshell’ signal region, are shown separately. The High and Intermediate p_T^{miss} regions dominate, with contributions primarily from $\tilde{\chi}_3^0 - \tilde{\chi}_2^\pm$ and $\tilde{\chi}_2^\pm - \tilde{\chi}_2^\pm$ pair-production – mechanisms which were absent in the original simplified model interpretation. This is similar to the exclusion seen in the compressed region for FullHad in Section 4.4.1. In this case, two leptons can be produced via the $\tilde{\chi}_3^0 \rightarrow Z\tilde{\chi}_1^\pm$ or $\tilde{\chi}_2^\pm \rightarrow Z\tilde{\chi}_1^\pm$ decays and jets via the $\tilde{\chi}_2^\pm \rightarrow W\tilde{\chi}_1^0$ decay. The compressed $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ then behave as a collective p_T^{miss} system.

The exclusion drops off towards the edges of the contour in Fig 4.20a due to deviations from the simplified model assumptions. Figure 4.21b shows the distribution of the $\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0$ branching ratio for models in the EWKino scan which lie inside the simplified model exclusion contour, before and after applying the 2L2J search. As expected, the exclusion is strongly correlated with this branching ratio. At the edge of the area of 2L2J sensitivity, where acceptance is smallest, models need to be closer to the simplified model (ie. have $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0 \approx 100\%)$ to be excluded.




Figure 4.20.: Fraction of EWKino scan models excluded by 2L2J in several mass planes. Grey bins contain no models and black bins have 0% exclusion. (a) is overlain with the observed exclusion contour for the simplified model in Figure 4.19.




Figure 4.21.: (a) Total inclusive truth-smearred 2L2J analysis signal region yield for each SUSY production process separated by signal region, for all models which are excluded by 2L2J (at truth-level) and lie outside the simplified model exclusion contour in the $\tilde{\chi}_2^0 - \tilde{\chi}_1^0$ mass plane with $m(\tilde{\chi}_1^\pm) < 300$ GeV. (b) Distribution of the $\tilde{\chi}_2^0 \rightarrow Z \tilde{\chi}_1^0$ branching ratio for models in the EWKino scan which lie inside the 2L2J expected exclusion contour. The ratio panel shows the fraction of models excluded by 2L2J.

Higgsino WZ (Figure 4.22a)	Wino/Bino WZ (Figure 4.22a)	Wino/Bino Wh (Figure 4.22b)
$(\tilde{\chi}_1^0, \tilde{\chi}_1^\pm, N2)$ are a higgsino triplet	$\tilde{\chi}_1^0$ is pure bino, $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ are pure wino	$\tilde{\chi}_1^0$ is pure bino
$\text{BR}(\tilde{\chi}_1^\pm \rightarrow W^{(*)} \tilde{\chi}_1^0) = 100\%$		
$\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^* \tilde{\chi}_1^0) = 100\%$	$\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^{(*)} \tilde{\chi}_1^0) = 100\%$	$\text{BR}(\tilde{\chi}_2^0 \rightarrow h \tilde{\chi}_1^0) = 100\%$
$m(\tilde{\chi}_1^\pm) = (m(\tilde{\chi}_2^0) + m(\tilde{\chi}_1^0))/2$	$m(\tilde{\chi}_1^\pm) = m(\tilde{\chi}_2^0)$	$m(\tilde{\chi}_1^\pm) = m(\tilde{\chi}_2^0)$

Table 4.11.: Simplified model assumptions corresponding to the models for the 3L analysis in Figure 4.22.

3L

The ‘3L’ analysis [144] targeted $\tilde{\chi}_1^\pm - \tilde{\chi}_2^0$ pair-production with a three-lepton final state. The signal regions were optimized for, and results interpreted in, several simplified SUSY models which are summarised in Table 4.11. Relevant diagrams are shown in Figure 4.22.




Figure 4.22.: Diagrams of the signatures targeted by the 3L analysis.

The analysis includes signal regions targeting decays via both onshell and offshell W , Z and Higgs bosons through a selection on the invariant mass of a same-flavour opposite sign lepton pair. The onshell signal regions are binned in m_T , while the offshell regions use selections on a variety of kinematic variables, including m_T , m_{T2} , $p_T^{\ell_{1,2,3}}$, and p_T^{miss} , to discriminate signal and background.

Figure 4.23 shows the exclusion from the 3L onshell (WZ) signal regions of the EWKino scan models in the $m(\tilde{\chi}_1^\pm) - m(\tilde{\chi}_1^0)$ and $m(\tilde{\chi}_2^0) - \Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ planes. Figure 4.23a includes bino-like LSP models from the EWKino scan and Figure 4.23b includes higgsino-like LSP models from the EWKino scan. The observed exclusion contour for the relevant simplified model in each case is overlain. The exclusion seen is consistent with the contours but weakens towards the edges. The distribution of $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0)$ for models in the EWKino scan that lie inside the 3L onshell (WZ) exclusion contour is shown in Figure 4.24. As expected, there is a clear correlation between this branching ratio and the fraction of models excluded.

Figure 4.25 shows the exclusion from the 3L (offshell) signal regions in the $m(\tilde{\chi}_2^0) - \Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane for (a) Bino-DM scan models and (b) higgsino-LSP EWKino scan models. The observed exclusion contours for the relevant simplified models are overlain. The bino-LSP models in Figure 4.25a are well covered at low $m(\tilde{\chi}_2^0)$, though the exclusion drops off towards higher $m(\tilde{\chi}_2^0)$ where the production cross-section and acceptance decreases. The region covered by the higgsino contour in Figure 4.25b is not well sampled by the higgsino models in the EWKino scan, which in general have smaller mass splittings. Nonetheless, the exclusion seen is consistent with the expected contour. The bulk of the




Figure 4.23.: Fraction of models excluded by the 3L onshell WZ signal regions in (a) the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane for bino-LSP models from the EWKino scan and (b) the $m(\tilde{\chi}_1^\pm) - \Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane for higgsino-LSP models from the EWKino scan. Grey bins contain no models and black bins have 0% exclusion. The observed exclusion contour for the relevant simplified model is overlain.




Figure 4.24.: Distribution of the $\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0$ branching ratio for EWKino scan models which lie inside the 3L onshell (WZ) simplified model exclusion contour. The ratio plot shows the fraction of models excluded by 3L onshell (WZ).




Figure 4.25.: Fraction of models excluded by the 3L (offshell) analysis in the 2D plane of $m(\tilde{\chi}_2^0)$ vs. $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$. Grey bins contain no models and black bins have 0% exclusion. The observed exclusion contour is overlain for the relevant simplified model in Table 4.11 for (a) bino-LSP models from the Bino-DM scan and (b) higgsino-LSP models from the EWKino scan.

higgsino-LSP models in the EWKino scan lie below the 3L (offshell) contour and are instead covered by the Compressed analysis.

4L

The search for SUSY in events with four or more leptons [145], referred to as 4L here, targeted a variety of models which can produce high lepton-multiplicity final states. Specifically, the analysis results were interpreted in R-parity violating scenarios in which the lightest neutralino is not stable and general gauge-mediated (GGM) SUSY models with a gravitino LSP. The latter is shown in Figure 4.26. These scenarios are distinct from the typical wino/higgsino production scenarios with a neutralino LSP targeted by the other searches described in this section. Therefore it is not possible to make direct comparisons between the 4L exclusion of pMSSM models and the simplified model results. Nonetheless, the fractional exclusion of the EWKino scan models projected into the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane is shown in Figure 4.27a.

The 4L analysis is expected to be sensitive to the production of the heavier charginos and neutralinos ($\tilde{\chi}_2^\pm, \tilde{\chi}_3^0, \tilde{\chi}_4^0$) followed by long decay chains that produce four or more leptons via W , Z and Higgs bosons. This can provide unique sensitivity to a range




Figure 4.26.: Diagrams of the signatures targeted by the 4L analysis.




Figure 4.27.: (a) Fraction of EWKino scan models excluded by 4L in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane. Grey bins contain no models and black bins have 0% exclusion. (b) Total inclusive truth-smear 4L analysis signal region yield for each SUSY production process, summed over all models in the EWKino scan that 4L is sensitive to (at truth-level).

of unconventional SUSY scenarios not covered by the other analyses. Indeed, Figure 4.27b shows the total (truth-smear) signal region yields for each SUSY production process, split by signal region lepton multiplicity and summed over the models in the EWKino scan that 4L is sensitive to (at truth-level). As expected, production of the heavier chargino and neutralinos dominates.

Disappearing track

The disappearing track analysis targeted compressed scenarios with long-lived charginos. The targeted model is shown in Figure 4.28. It includes a wino or higgsino $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_1^0$. The $\tilde{\chi}_1^\pm$ is long-lived, and leaves a measurable track in the inner detector before decaying to $\tilde{\chi}_1^0$ and a soft pion which is not reconstructed – resulting in a “disappearing” track.

As outlined in Section 4.2.3, this analysis is not evaluated using SIMPLEANALYSIS or RECAST as the event generation used is not set up to correctly handle long-lived particles. Instead, the upper limits on production cross-sections are used to determine exclusion. Figure 4.29a shows a 2D scatter plot of the pMSSM models from both the EWKino and Bino-DM scans that are excluded in the 2D plane of the chargino mass vs. lifetime. Model points in orange are outside the region for which a valid interpolation of the available cross-section limits can be made. Figure 4.29b shows the fraction of models from the Bino-DM and EWKino scans excluded by disappearing track in the $m(\tilde{\chi}_1^\pm) - \Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ plane. Disappearing track is the only analysis which directly targets scenarios with a wino-like LSP, where the mass splitting between the lightest chargino and neutralino is less than 1 GeV. It therefore provides unique sensitivity in the highly compressed region, though as seen previously the FullHad and 2L2J analyses are also sensitive to some of these scenarios via decays of heavier electroweakinos.




Figure 4.28.: Diagram of the signature targeted by the Disappearing Track analysis.




Figure 4.29.: (a) Scatter of pMSSM models in the chargino mass–lifetime plane. Both the EWKino and Bino-DM scans are included. Models coloured in blue are excluded by the disappearing track cross-section upper limits, in green are not excluded and in orange are outside the valid interpolation region. (b) Shows the fraction of EWKino and Bino-DM scan models excluded by the disappearing track analysis in the $m(\tilde{\chi}_1^\pm) - \Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ plane. Black bins have 0% exclusion.

4.4.2. ATLAS bSM Higgs constraints

Two additional ATLAS results which constrain bSM Higgs physics are also used to exclude the pMSSM models.

First, in Reference [148] a combination of searches for invisible decays of the Higgs boson is performed, placing upper limits on $\text{BR}(h \rightarrow \text{invisible})$. In the pMSSM, the invisible decay mode of the Higgs boson is $h \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0$, and the upper limits can be used to exclude models with $\text{BR}(h \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0)$ above the observed limit. Figures 4.30a and 4.30b show the distribution of $\text{BR}(h \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0)$ before and after applying the direct SUSY searches described so far for the EWKino and Bino-DM scans respectively. The expected and observed upper limits from Reference [148] (0.077 and 0.107 respectively) are displayed in red. Note that the Bino-DM scan $\text{BR}(h \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0)$ distribution has an enhanced tail which contains models in the Higgs funnel region. The $\text{BR}(h \rightarrow \text{invisible})$ limit excludes only a few models that are not already excluded by the direct searches.

Second, in Reference [147] a combined Higgs boson cross-section and branching ratio measurement is used to place limits on the mass of the pseudoscalar Higgs boson, $m(A)$. Figures 4.30c and 4.30d show the distribution of $m(A)$ before and after applying the direct SUSY searches described so far for the EWKino and Bino-DM scans respectively. The observed lower limit from Reference [147], $m(A) > 480$ GeV, is displayed in red. The difference in the shape of this distribution between the two scans is due to the preference for $m(A) \approx 2 \times m(\tilde{\chi}_1^0)$ in the Bino-DM scan so that the A -funnel annihilation mechanism $\tilde{\chi}_1^0 \tilde{\chi}_1^0 \rightarrow A$ can reduce the dark matter relic density.

4.4.3. Overlaps and unique sensitivity

The overlap in sensitivity between the analyses discussed in this section is displayed in Figure 4.31. The number in each box indicate the percentage of the models excluded by the analysis on the y -axis which are also excluded by the analysis on the x -axis. For example in Figure 4.31a, of the models excluded by FullHad, 22% are also excluded by 2L2J (top row, fifth column). For most analyses, the largest overlap is with the 2L2J and 3L analyses which covered the largest area of the $\tilde{\chi}_1^\pm / \tilde{\chi}_2^0 - \tilde{\chi}_1^0$ mass planes. Disappearing Track provides mostly unique exclusion, though there is some overlap with FullHad and 2L2J as discussed previously.

Figure 4.32 indicates models with unique (expected) sensitivity for each search. All excluded models are shown (except for Figure 4.32c where only excluded models with unique analysis sensitivity are shown), and each model point is coloured according to the analysis with unique sensitivity (or coloured silver if multiple analyses have sensitivity). In general, each analysis displays some unique sensitivity in the region corresponding to its reported simplified model exclusion. FullHad and 2L2J notably provide unique sensitivity to models with small mass-splittings in the EWKino scan in Figure 4.32b, as discussed previously. Also notable is the unique sensitivity of 1Lbb in Figure 4.31a. This search had a relatively poor fractional exclusion in the $\tilde{\chi}_2^0 - \tilde{\chi}_1^0$ mass plane. However, it provides a significant amount of unique exclusion as the only search that specifically targeted the $\tilde{\chi}_2^0$ decay via a Higgs (rather than Z) boson.




Figure 4.30.: Distribution of $\text{BR}(h \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0)$ (first row) and $m(A)$ (second row) before and after the ATLAS Run 2 direct searches. The first column shows the EWKino scan and the second column the Bino-DM scan. Relevant upper/lower limits are shown as red lines. Only models with $m(\tilde{\chi}_1^\pm) < 1200$ GeV and $m(\tilde{\chi}_1^0) < 500$ GeV are included to focus on the region the ATLAS searches are sensitive to.




Figure 4.31.: Percentage overlaps in the exclusion of each ATLAS analysis. For each square, the number and colour indicates the percentage of the models excluded by the analysis on the y -axis which are also excluded by the analysis on the x -axis, ie. $100 \times N_{xy}/N_y$ where N_{xy} is the number of models excluded by the results on the x AND y axes and N_y is the number of models excluded by the result on the y axis.




Figure 4.32.: Scatter plots coloured by the search (not including $m(A)$ and $\text{BR}(h \rightarrow \text{inv.})$ constraints) with unique (expected) sensitivity. Models that multiple searches are sensitive to are shown in silver. Only models that are excluded are shown and for clarity in (c) only models with unique analysis sensitivity are shown due to the density of points. The first row shows the EWKino scan and the second row the Bino-DM scan. The first column shows the $m(\tilde{\chi}_1^\pm) - m(\tilde{\chi}_1^0)$ plane and the second column the $m(\tilde{\chi}_2^0) - \Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane.

4.5. External constraints

In addition to the ATLAS results discussed in the previous section, various non-ATLAS ‘external’ constraints can be used to constrain the pMSSM models. The external constraints considered in this analysis are summarised in Table 4.12. These constraints are applied to the EWKino and Bino-DM model sets before producing all results in the following sections. In this section, each constraint will be discussed and its impact on the pMSSM scan(s) presented. Note that constraints related to dark matter are not included in Table 4.12 and will be discussed separately in Section 4.7. The EWKino scan and Bino-DM scan have similar distributions and selection efficiencies for each constraint variable. Therefore for brevity, the distribution of each constraint variable will be shown for the EWKino scan only.

LEP

Searches for charginos at LEP [138] placed a lower bound of 91.9 GeV on the mass of the lightest chargino, and a stricter bound of 103 GeV for scenarios with moderate mass

External constraints			
Constraint	Lower bound	Upper bound	Reference
LEP			
$m(\tilde{\chi}_1^\pm)$	103 GeV ($\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) \geq 3$ GeV) 91.9 GeV ($\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) < 3$ GeV)		Combined LEP constraint [138]
Higgs mass			
$m(h)$	120 GeV	130 GeV	Window around $m(h) \approx 125$ GeV [173]
Flavour			
$\text{BR}(b \rightarrow s\gamma)$	3.11×10^{-4}	3.87×10^{-4}	2022 PDG average $\pm 2\sigma$ [13]
$\text{BR}(B_s \rightarrow \mu\mu)$	1.87×10^{-9}	4.31×10^{-9}	Most recent LHCb result $\pm 2\sigma$ [174]
$\text{BR}(B^+ \rightarrow \tau\nu)$	6.10×10^{-5}	1.57×10^{-4}	2022 PDG average $\pm 2\sigma$ [13]
EW precision			
$\Delta\rho$	-0.0004	0.0018	GFITTER group [175]
$\Gamma_{\text{inv}}(Z)$		2 MeV	SLC and LEP measurements [176].
$m(W)$	80.347 GeV	80.407 GeV	2022 PDG average $\pm 2\sigma$ [13] • excluding CDF measurement [177] • ± 6 MeV from u/c on $m(t)$ [178]


Table 4.12.: Additional non-ATLAS constraints from previous collider searches, Higgs measurements, flavour physics observables and electroweak precision measurements.

splittings. As specified in Section 4.2.1, all models failing this constraint are removed from the EWKino and Bino-DM scans by default.

Higgs boson mass

The distribution of the Higgs boson mass is shown in Figure 4.33 for all EWKino scan models and those that survive ATLAS Run 2. The Higgs boson mass has been measured precisely – for example, Reference [173] measured $m(h) = 124.94 \pm 0.17 \pm 0.03$ GeV. However, applying such a tight constraint to the pMSSM models would be unnecessarily restrictive.

The calculated Higgs mass in the MSSM is influenced by the masses of the various sparticles, in particular the mass of the stop. In the pMSSM scans performed here,



(a)

Figure 4.33.: Distribution of the Higgs boson mass before and after ATLAS Run 2 for the EWKino scan model set. The red dashed lines indicate the 120 GeV–130 GeV mass window used to constrain the model set. The ratio panel shows the fraction of models excluded by ATLAS. Only models with $m(\tilde{\chi}_1^\pm) < 1200$ GeV and $m(\tilde{\chi}_1^0) < 500$ GeV are shown to focus on the region ATLAS searches are sensitive to.

the stop is decoupled with a mass parameter varying between 2–5 GeV and does not impact the sensitivity of the searches considered (event generation is not performed for the production of stops and other squarks). Therefore, for a given model the Higgs mass could be arbitrarily corrected to the experimental measured value by modifying the masses of the decoupled sparticles. However, a loose window of 120 GeV – 130 GeV is imposed to ensure that the kinematics of the Higgs boson in the generated events is “Standard Model-like” and similar to that assumed in the development of each ATLAS search. Within this window, the ATLAS sensitivity to the pMSSM models is found to be independent of the Higgs mass, as shown in the ratio panel of Figure 4.33. As specified in Section 4.2.1, this Higgs mass window is used to constrain the EWKino and Bino-DM model prior to producing all results shown in this chapter other than Figure 4.33.

Flavour constraints

The branching ratios $\text{BR}(b \rightarrow s\gamma)$, $\text{BR}(B_s \rightarrow \mu\mu)$ and $\text{BR}(B^+ \rightarrow \tau\nu)$ are all required to lie within two standard deviation (2σ) intervals around the latest experimental measurements. In the pMSSM models, loop contributions involving sparticles impact these branching ratios such that they may disagree with the experimental measurements. The distributions of each are shown in Figure 4.34 for all models in the EWKino scan and those that survive the ATLAS Run 2 results. In each plot, the constraint window used is indicated by red dashed lines. For $\text{BR}(b \rightarrow s\gamma)$ and $\text{BR}(B^+ \rightarrow \tau\nu)$ much of the tails in the distributions that lie outside the constraint window are also excluded by ATLAS – mainly by the ATLAS constraint $m(A) > 480$ GeV. The ATLAS exclusion is fairly independent of $\text{BR}(B_s \rightarrow \mu\mu)$ on the other hand.

EW precision constraints

The distributions of three electroweak precision observables are shown in Figure 4.34 for all models in the EWKino scan and those that survive the ATLAS Run 2 results. At the time of writing, some controversy exists around the measured mass of the W boson, with the latest measurement from CDF [177] conflicting with both the Standard Model prediction and other experimental measurements. In this analysis, the PDG average of experimental measurements of the W boson mass, *excluding* the CDF measurement in Reference [177], is used to constrain the pMSSM models. The CDF measurement is displayed separately in Figure 4.35a. If confirmed it would provide evidence of bSM




Figure 4.34.: Distribution of flavour observables for EWKino scan models before and after ATLAS Run 2. The red dashed lines indicate the window used when applying each constraint to the model sets. The ratio panel shows the fraction of models excluded by ATLAS. Only models with $m(\tilde{\chi}_1^\pm) < 1200 \text{ GeV}$ and $m(\tilde{\chi}_1^0) < 500 \text{ GeV}$ are shown to focus on the region ATLAS searches are sensitive to.




Figure 4.35.: Distribution of electroweak precision observables for EWKino scan models before and after ATLAS Run 2. The red dashed lines indicate the bound(s) used when applying each constraint to the model sets. The ratio panel shows the fraction of models excluded by ATLAS. Only models with $m(\tilde{\chi}_1^\pm) < 1200$ GeV and $m(\tilde{\chi}_1^0) < 500$ GeV are shown to focus on the region ATLAS searches are sensitive to.

physics and point to a narrow range of pMSSM scenarios – 3.6% of EWKino scan models lie in the displayed CDF range. The W mass window used to constrain the models (and displayed in Figure 4.35a) is expanded by 6 MeV to account for the experimental uncertainty of the top-quark mass which impacts the MSSM W mass calculation [178].

The electroweak parameter $\Delta\rho$, displayed in Figure 4.35c, probes the radiative corrections to the Z boson coupling strength, the effective weak mixing angle, and the W boson mass. It is constrained by the global fit result on the oblique parameter [179, 180] $T = 0.09 \pm 0.14$ in Reference [175]. The $\Delta\rho$ parameter is given by $\Delta\rho = \alpha T$ where $\alpha = 1/128$. The $Z \rightarrow$ invisible decay width, shown in Figure 4.35b, is also constrained by measurements from LEP and SLC to be less than 2 MeV [176]. However, all models in the EWKino scan (that pass the LEP and Higgs mass constraints) pass this constraint.

Overlaps

The overlap in sensitivity between the constraints discussed in this section, and the ATLAS results from the previous section, is displayed in Figure 4.36. The number in each box indicates the percentage of the models excluded by the constraint on the y -axis which are also excluded by the constraint on the x -axis.




Figure 4.36.: Percentage overlaps in the exclusion of each constraint. For each square, the number and colour indicate the percentage of the models excluded by the result on the y -axis which are also excluded by the result on the x -axis, ie. $100 \times N_{xy}/N_y$ where N_{xy} is the number of models excluded by the results on the x AND y axes and N_y is the number of models excluded by the result on the y axis.

4.6. Overall ATLAS constraints on electroweakinos

In this section, the overall ATLAS exclusion of the EWKino and Bino-DM scans is presented. To determine whether a model is ‘excluded’, the analysis with the best (ie. smallest) expected $\text{CL}_s^{\text{Exp.}}$ of the ATLAS searches discussed in Section 4.4 is used. If the observed $\text{CL}_s^{\text{Obs.}}$ for that analysis is less than 0.05, the model is considered to be excluded. Models that do not satisfy the ATLAS $m(A)$ and $\text{BR}(h \rightarrow \text{inv.})$ constraints are also categorised as ‘excluded by ATLAS’. Only models which pass all external constraints in Table 4.12 are included here, such that the ATLAS exclusion of otherwise viable⁵ models is being studied.

Electroweakino masses

Figure 4.37 shows the fraction of models excluded by ATLAS in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane for the (a) EWKino and (b) Bino-DM scans. To indicate the region of simplified models excluded by ATLAS, the overlain contour shows the envelope of the 3L, 2L2J and FullHad analysis observed exclusion of a $\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W \tilde{\chi}_1^0 Z \tilde{\chi}_1^0$ simplified model with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^{(*)} \tilde{\chi}_1^0) = 100\%$.

In general, the regions excluded match the areas covered by the simplified model contours. However, it is notable that only a handful of bins show 100% exclusion, indicating that despite the optimistic simplified model exclusion contours, many scenarios remain uncovered, even at relatively low $\tilde{\chi}_1^\pm$ mass. As discussed for individual analyses in Section 4.4, the fraction of models excluded decreases towards the edges of the overlain contour, where analysis acceptance decreases such that there is no sensitivity to models with smaller branching fractions and cross-sections compared to simplified models. The distribution of models is significantly different between the EWKino scan and the Bino-DM scan. The EWKino scan fills the plane approximately uniformly, while the dark matter relic density requirement imposed on the Bino-DM model sample concentrates models in the funnel regions with $m(\tilde{\chi}_1^0) \approx m_{Z,h}/2$ and in the compressed regions.

Figure 4.38 shows the fraction of models excluded by ATLAS in the $\tilde{\chi}_2^0 - \tilde{\chi}_1^0$ mass plane for the (a) EWKino and (b) Bino-DM scans. The same contour as Figure 4.37 is overlain. The exclusion pattern is similar to the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane, but for the EWKino scan

⁵The models do not necessarily satisfy dark-matter-related constraints which are discussed separately in Section 4.7. All Bino-DM models do satisfy the dark matter relic density constraint by construction though.




Figure 4.37.: Fraction of models excluded by Run 2 ATLAS results in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane for the (a) EWKino and (b) Bino-DM scans. Only models which pass all non-dark-matter external constraints are included. Grey bins contain no models and black bins have 0% exclusion. The overlain contour shows the envelope of the FullHad + 2L2J + 3L exclusion of a $\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W \tilde{\chi}_1^0 Z \tilde{\chi}_1^0$ simplified model with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^{(*)} \tilde{\chi}_1^0) = 100\%$.




Figure 4.38.: Fraction of models excluded by Run 2 ATLAS results in the $\tilde{\chi}_2^0 - \tilde{\chi}_1^0$ mass plane for the (a) EWKino and (b) Bino-DM scans. Only models which pass all non-dark-matter external constraints are included. Grey bins contain no models and black bins have 0% exclusion. The overlain contour shows the envelope of the FullHad + 2L2J + 3L exclusion of a $\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W \tilde{\chi}_1^0 Z \tilde{\chi}_1^0$ simplified model with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^{(*)} \tilde{\chi}_1^0) = 100\%$.

sensitivity is seen beyond the overlain contour for $\tilde{\chi}_2^0$ masses up to 2 TeV. These are scenarios with a wino-like $\tilde{\chi}_1^0$ and $\tilde{\chi}_1^\pm$ and bino- or higgsino-like $\tilde{\chi}_2^0$ with a large mass splitting. The disappearing track analysis is sensitive to the highly compressed $\tilde{\chi}_1^0$ and $\tilde{\chi}_1^\pm$ in these scenarios.

Exclusion outside of the overlain contours can be seen in the compressed regions of Figures 4.37 and 4.38. These regions can be seen explicitly in Figures 4.39 and 4.40 which show the fraction of models excluded by ATLAS in the $m(\tilde{\chi}_2^0)$ vs. $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ and $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ planes respectively.

Figure 4.39a shows models from the EWKino scan and is overlain with the envelope of the 3L and Compressed analysis observed exclusion of the ATLAS higgsino simplified model. Inside the contour, most bins display over 50% exclusion, and the exclusion also extends beyond the region covered by the contour. The FullHad and 2L2J analyses have sensitivity to many of these compressed models, as discussed in Section 4.4. Production mechanisms involving higher mass electroweakinos such as $\tilde{\chi}_2^\pm$ and $\tilde{\chi}_4^0$ are contributing here with the compressed $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ acting as a collective p_T^{miss} system. Indeed, the sensitivity halts at around $\tilde{\chi}_2^0 \approx 450$ GeV, which corresponds to the upper limit on $\tilde{\chi}_1^0$ for simplified models excluded by FullHad and 2L2J.

Figure 4.39b shows models from the Bino-DM scan and is overlain with the envelope of the 3L and Compressed analysis observed exclusion of a wino- $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ and bino-LSP simplified model. Strong exclusion is seen for the lowest mass bins, but this reduces towards the edge of the overlain contour. Exclusion is seen beyond the overlain contour for small mass splittings ($\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) < 5$ GeV) up to $m(\tilde{\chi}_2^0) \approx 500$ GeV. This is mainly due to the disappearing track analysis which is not included in the overlain contour.

Figure 4.40 shows the fraction of models excluded by ATLAS in the $m(\tilde{\chi}_1^\pm)$ vs. $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ plane. The EWKino scan is shown in Figure 4.40a for highly compressed scenarios with $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) < 1$ GeV. These are predominantly wino-like LSP models and sensitivity is provided by the disappearing track analysis and FullHad/2L2J from decays of heavier electroweakinos as discussed previously. For the Bino-DM scan, shown in Figure 4.40b, the models in this compressed region mainly contain a wino-like $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ which are close in mass to the bino-like LSP to reduce the dark matter relic density through $\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ co-annihilation. The disappearing track analysis is providing the bulk of the sensitivity here.

Figures 4.41a and 4.41b show the distribution of the $\tilde{\chi}_1^0$ mass for the EWKino and Bino-DM scans respectively. The distributions are shown separately for a bino-, wino-




Figure 4.39.: Fraction of models excluded by Run 2 ATLAS results in the $m(\tilde{\chi}_2^0)$ vs. $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane for the (a) EWKino and (b) Bino-DM scans. Only models which pass all non-dark-matter external constraints are included. Grey bins contain no models and black bins have 0% exclusion. The overlain contours show the ATLAS 3L + Compressed exclusion of a $\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W^* \tilde{\chi}_1^0 Z^* \tilde{\chi}_1^0$ simplified model with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^{(*)} \tilde{\chi}_1^0) = 100\%$ and (a) a pure higgsino LSP and $\tilde{\chi}_2^0$, (b) a pure bino LSP and pure wino $\tilde{\chi}_2^0$.




Figure 4.40.: Fraction of models excluded by Run 2 ATLAS results in the $m(\tilde{\chi}_1^\pm)$ vs. $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ plane for the (a) EWKino and (b) Bino-DM scans. Only models which pass all non-dark-matter external constraints are included. Grey bins contain no models and black bins have 0% exclusion.

and higgsino-like $\tilde{\chi}_1^0$. In the Bino-DM scan, an LSP with a mass below 100 GeV is almost entirely excluded, but the sensitivity rapidly drops off with increasing mass. Most models with an LSP mass greater than 200 GeV remain viable. In the EWKino scan, Bino-LSPs of all masses remain viable, with only up to around 50% of models excluded in any bin. A wino-like LSP is the most constrained by the ATLAS results, with greater than 50% exclusion up to around $m(\tilde{\chi}_1^0) \approx 450$ GeV.

Figures 4.41c and 4.41d show the distribution of the $\tilde{\chi}_1^\pm$ mass for the EWKino and Bino scans respectively. The distributions are shown separately for a wino- and higgsino-like $\tilde{\chi}_1^\pm$. The exclusion fractions shown in the ratio panels differ for each scan. In the EWKino scan, a larger exclusion fraction is seen for a wino-like $\tilde{\chi}_1^\pm$ while in the Bino-DM scan a larger exclusion fraction is seen for a higgsino-like $\tilde{\chi}_1^\pm$. This relates to the typical spectrum in each scan as follows.

In the EWKino scan, a higgsino-like $\tilde{\chi}_1^\pm$ is typically part of a compressed-mass triplet with $\tilde{\chi}_1^0$ and $\tilde{\chi}_2^0$ (with mass splittings around a few GeV). The Compressed and 3L (offshell) analyses provide some sensitivity to these scenarios. A wino-like $\tilde{\chi}_1^\pm$ in the EWKino scan is either very close in mass to a wino-like $\tilde{\chi}_1^0$ or well separated from the $\tilde{\chi}_1^0$ and instead nearly mass-degenerate with $\tilde{\chi}_2^0$. In the former case, the disappearing track analysis provides good coverage, while the latter case was targeted by the 1Lbb, 2L0J, 2L2J and 3L searches. Therefore overall more exclusion is seen for wino-like rather than higgsino-like $\tilde{\chi}_1^\pm$'s.

In the Bino-DM scan, however, a model with a higgsino-like $\tilde{\chi}_1^\pm$ is most likely a funnel-region model. Here the LSP is bino-like with a moderate mass splitting from the higgsino-like $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_3^0$ such that the direct non-compressed searches have good sensitivity. A model with a wino-like $\tilde{\chi}_1^\pm$ in the Bino-DM scan will most likely contain a bino-like LSP with a small mass-splitting to the wino-like $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ such that co-annihilation is enhanced to reduce the relic density. These compressed scenarios primarily receive coverage from the disappearing track analysis, but the other direct searches have limited sensitivity. Therefore overall more exclusion is seen for higgsino-like rather than wino-like $\tilde{\chi}_1^\pm$'s.

Finally, Figure 4.42 shows a summary of the fractional exclusion of each electroweakino mass in the (a) EWKino and (b) Bino-DM scans. The majority of scenarios remain uncovered by current searches, with 100% exclusion only seen in the very lowest mass bins. It is emphasised that one must be careful when interpreting the fractional exclusion shown here and in other Figures. These fractions are highly dependent on the model




Figure 4.41.: Distribution of the LSP mass (first row) and $\tilde{\chi}_1^\pm$ mass (second row) before and after ATLAS Run 2, split by bino-/wino-/higgsino-type (note that in (b) all LSPs are bino-like by construction). The first column shows models in the EWKino scan that pass all non-dark-matter constraints and the second column models in the Bino-DM scan that pass all non-dark-matter constraints. The ratio panels show the fraction of models excluded by ATLAS in each case.




Figure 4.42.: Summary plots showing the fraction of models excluded by ATLAS across the masses of each electroweakino for models which pass non-dark-matter external constraints. (a) shows the EWKino scan and (b) the Bino-DM scan. Grey bins contain no models and black bins have 0% exclusion.

sample used – this can be seen by comparing the EWKino and Bino-DM scan for example. Therefore a qualitative rather than quantitative message should be drawn from these results. However, even relatively low-mass scenarios that one may assume to be excluded based on published simplified model limits remain viable. Overall, electroweakinos remain vastly under-explored after Run 2.

Branching ratios

Figure 4.43 shows the distributions of $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z)$, $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h)$ and $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^\pm W^\mp)$ before and after ATLAS Run 2. The first column shows the EWKino scan and the second column the Bino-DM scan. Note that these are branching fractions for two-body decays via onshell Z , h or W bosons. Therefore the large peaks at zero in all of these distributions contain models with compressed spectra for which three-body offshell decay modes dominate.

For the EWKino scan, the plots are separated by LSP type as the distribution and ATLAS sensitivity is distinct for each. These onshell decay modes are not relevant to models with a higgsino-like LSP so these are not shown. For models with a bino-like LSP in both the EWKino and Bino-DM scan, the highest exclusion by the ATLAS results occurs for large values of $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z)$, as expected due to the large number




Figure 4.43.: $\tilde{\chi}_2^0$ branching ratio distributions for the EWKino scan (first column) and Bino-DM scan (second column) before and after ATLAS Run 2. The EWKino scan distributions are split by bino/wino LSP type. Only models that pass all non-dark-matter constraints are included. The ratio panels show the fraction of models excluded by ATLAS. Note that $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^\pm W^\mp)$ refers to the sum of the W^+ and W^- modes.

of searches optimised for simplified models with $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z) = 100\%$ whilst only the 1Lbb search directly targets the $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h$ decay for low mass $\tilde{\chi}_2^0$. This gap will likely be filled by upcoming ATLAS searches for $\tilde{\chi}_2^0$'s that decay via a Higgs boson. For models with a wino-like LSP, the largest exclusion occurs when $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z) \approx 15\%$ and larger $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h)$. These branching ratios are impacted by the mass splitting between the $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$, which also affects the lifetime of the $\tilde{\chi}_1^\pm$. This peak in exclusion corresponds to models with a chargino lifetime of around 0.1 ns which aligns with the region of sensitivity of the Disappearing Track analysis.

4.7. Complementarity with dark matter direct-detection experiments

Direct dark matter detection experiments can also provide sensitivity to SUSY scenarios that include a dark matter candidate. These search explicitly for interactions between dark matter (in our case the LSP) and protons/neutrons through scattering processes like $p\tilde{\chi}_1^0 \rightarrow p\tilde{\chi}_1^0$. This is complementary to the collider searches considered so far that target the production of heavier electroweakinos which decay to the LSP via processes like $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W\tilde{\chi}_1^0 Z\tilde{\chi}_1^0$. Limits from two direct detection experiments are considered in this section:

1. LUX-ZEPLIN (LZ) [181] – at the time of writing this provides the most stringent limits on the spin-independent WIMP-nucleon scattering cross-section.
2. PICO-60 [182] – at the time of writing this provides the most stringent limits on the spin-dependent WIMP-proton scattering cross-section.

In this section, the sensitivity of these experiments and complementarity with ATLAS is presented for bino-LSP scenarios in the Bino-DM scan and wino-/higgsino-LSP scenarios in the EWKino scan separately. Only models that pass all other non-dark-matter external constraints are included.

Note that the limits published by LZ and PICO-60 are for a dark matter candidate that saturates the observed dark matter relic density, $\Omega h^2 = 0.12$. The majority of models in the pMSSM scans here instead produce less than the observed relic density. Therefore for a fair comparison with the limits, all WIMP-nucleon scattering cross-sections shown in this section are scaled by $(\Omega h^2)/0.12$ (ie. the ratio of the expected to observed relic

density) for each model. This obtains the effective cross-sections assuming the remaining non-LSP dark matter is invisible to the direct detection experiments.

4.7.1. Bino-like LSP

For the Bino-DM scan, the complementarity between ATLAS and the LZ and PICO-60 direct detection experiments can be seen in Figure 4.44. Figures 4.44a and 4.44b show the LSP mass versus the spin-independent WIMP-nucleon scattering cross-section overlain with the LZ limit. Figures 4.44c and 4.44d show the LSP mass versus the spin-dependent WIMP-proton scattering cross-section overlain with the PICO-60 limit. Models from the Bino-DM scan are shown both before and after the Run 2 ATLAS results, coloured by the dominant dark matter annihilation mechanism. Similarly, Figure 4.45 shows $m(\tilde{\chi}_1^0)$ versus Ωh^2 for (a) models in the Bino-DM scan, (b) those that survive the PICO-60 and LZ limits, (c) those not excluded by ATLAS, and (d) that that survive the PICO-60 and LZ limits AND are not excluded by ATLAS.

The direct detection experiments strongly constrain the Z/h funnel region in particular. ATLAS provides good complementarity and excludes the majority of Z/h funnel region models that lie below the direct detection limits. Overall, just a handful of Z/h funnel region scenarios remain viable after both ATLAS Run 2 and the direct detection experiments.

In the bulk region, almost all scenarios with dominant self-annihilation mechanisms to SM particles, such as $\tilde{\chi}_1^0 \tilde{\chi}_1^0 \rightarrow t\bar{t}$, are excluded by direct detection experiments. Models that have $\tilde{\chi}_1^\pm / \tilde{\chi}_2^0$ co-annihilation as the dominant mode have smaller WIMP-nucleon cross-sections and mostly lie below the direct detection limits. ATLAS does reduce the density of these scenarios but most, particularly above $m(\tilde{\chi}_1^0) \approx 200$ GeV, are still viable. These are mostly scenarios with a nearly pure-wino $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ close in mass to a pure-bino LSP. Additionally, models in the A/H funnel are also viable. These can have a larger LSP mass without requiring a compressed mass-spectrum.

4.7.2. Wino- and higgsino-like LSP

Figure 4.46 shows, for models in the EWKino scan with a wino- or higgsino-like LSP, the distribution of (a) the WIMP-nucleon spin-independent scattering cross-section and (b) the WIMP-proton spin-dependent scattering cross-section. The distribution is




Figure 4.44.: Models from the Bino-DM scan which pass all non-dark-matter external constraints, coloured by the dominant annihilation mechanism, before (first column) and after (second column) ATLAS Run 2. (a) and (b) show the $m(\tilde{\chi}_1^0)$ vs. WIMP-nucleon spin-independent scattering cross-section plane overlain with the limit from LUX-ZEPLIN [181]. (c) and (d) show the $m(\tilde{\chi}_1^0)$ vs. WIMP-proton spin-dependent scattering cross-section plane overlain with the limit from PICO-60 [182]. The pMSSM model cross-sections have been scaled by $(\Omega h^2)/0.12$ for a fair comparison with the limits.




Figure 4.45.: Models from the Bino-DM scan which pass all non-dark-matter external constraints in the $m(\tilde{\chi}_1^0)$ vs. Ωh^2 plane, coloured by the dominant annihilation mechanism. (a) includes all models, (b) models surviving constraints from direct detection experiments, (c) models surviving ATLAS Run 2, and (d) models surviving both ATLAS Run 2 and constraints from direct detection experiments.




Figure 4.46.: Distributions of (a) the WIMP-nucleon spin-independent scattering cross-section before and after LUX-ZEPLIN [181] and ATLAS Run 2, and (b) the WIMP-proton spin-dependent scattering cross-section before and after PICO-60 [182] and ATLAS Run 2. Models from the EWKino scan with a higgsino- or wino-like LSP that pass all non-dark-matter constraints are included. The pMSSM model cross-sections have been scaled by $(\Omega h^2)/0.12$ for a fair comparison with cross-section limits from direct detection experiments.

shown for models that pass all non-dark-matter external constraints, after limits from the relevant direct detection experiment have been applied, and after constraints from ATLAS Run 2 have additionally been applied. This reveals the complementarity between direct detection experiments and ATLAS searches for wino- and higgsino-LSP scenarios. ATLAS provides good exclusion of low cross-section scenarios which lie well below the LZ and PICO-60 limits.

Finally, Figure 4.47 shows $m(\tilde{\chi}_1^0)$ versus Ωh^2 for (a) all models with a wino- or higgsino-like LSP in the EWKino scan, (b) those that survive the PICO-60 and LZ limits, (c) those not excluded by ATLAS, and (d) that survive the PICO-60 and LZ limits AND are not excluded by ATLAS. The points are coloured with RGB value set to the wino/higgsino/bino fraction of the LSP and the dashed horizontal line indicates the observed dark matter relic density. Scenarios with a mixed LSP are mostly excluded by the direct detection experiments and ATLAS excludes most low-mass wino scenarios. However higgsino-LSP scenarios remain viable at masses as low as 91 GeV, which is on the boundary of the LEP limit. Additionally, wino-LSP scenarios that saturate the




Figure 4.47.: Models from the EWKino scan with a wino- or higgsino-like LSP that pass all non-dark-matter external constraints in the $m(\tilde{\chi}_1^0)$ vs. Ωh^2 plane, coloured with RGB value set to the wino/higgsino/bino fraction of the LSP. (a) includes all models, (b) models surviving constraints from direct detection experiments, (c) models surviving ATLAS Run 2, and (d) models surviving both ATLAS Run 2 and constraints from direct detection experiments. The dashed horizontal line indicates the observed dark matter relic density.

observed relic density are excluded by direct detection experiments, while higgsino-LSP models with $\Omega h^2 = 0.12$ remain viable. However, these scenarios have an LSP mass around 1 TeV which is well out of reach for any ATLAS searches in the near future.

4.8. Summary & outlook

4.8.1. Summary of exclusion

Table 4.13 contains a summary of the percentage of models excluded for each LSP type, for each search and constraint individually and overall. In total, the results presented in this section can be viewed through two lenses. On the one hand, the final row of Table 4.13 shows that over 80% of models for each LSP type and scan are excluded when considering both ATLAS Run 2 and external constraints. This could be interpreted as disfavouring the viability of SUSY. However, as has been noted previously, a quantitative interpretation is not entirely appropriate given the dependence of the exclusion fraction on the priors, ranges and sampling strategy of the scan. So on the other hand, reassurance could be taken from the results presented in the previous section, which showed non-excluded models with masses of a few hundred GeV and less – well within the reach of Run 3 at the LHC.

4.8.2. Viable benchmark scenarios

It is important to understand why these scenarios have been missed by previous searches such that they can be targeted with future searches. Several benchmark scenarios are presented that pass all constraints and have $\text{LSP}/\tilde{\chi}_1^\pm/\tilde{\chi}_2^0$ masses which lie within published simplified model exclusion contours. Only benchmark scenarios with a bino- or higgsino-like LSP are discussed. Since only the disappearing track search directly targeted wino-like LSP scenarios, and is applied as cross-section upper limits, the reason for non-excluded wino-LSP scenarios is simpler: the production cross-section is not large enough.

Search	EWKino scan			Bino-DM scan
	Bino-LSP	Higgsino-LSP	Wino-LSP	
FullHad	5.28 %	2.57 %	2.97 %	0.24 %
1Lbb	1.23 %	0.00 %	0.00 %	0.10 %
Compressed	0.23 %	1.35 %	0.00 %	1.91 %
2L0J	2.19 %	0.00 %	0.04 %	0.74 %
2L2J	5.55 %	2.89 %	3.11 %	15.77 %
3LOffshell	1.41 %	0.61 %	0.00 %	11.68 %
3L Onshell WZ	4.07 %	0.74 %	1.18 %	13.29 %
3L Onshell Wh	0.92 %	0.09 %	0.00 %	6.13 %
4L	0.31 %	0.09 %	0.04 %	2.66 %
Disappearing Track	0.06 %	0.13 %	45.7 %	1.87 %
Overall ATLAS Run 2 searches	11.24 %	5.70 %	48.45 %	24.88%
$m(A) > 480$ GeV	0.43 %	1.63 %	0.43 %	4.77 %
$\text{BR}(h \rightarrow \text{inv.}) \leq 0.107$	0.37 %	0.00 %	0.00 %	8.03 %
Overall ATLAS Run 2	11.61 %	7.22 %	48.70 %	28.13%
Flavour constraints	25.66 %	28.87 %	25.72 %	34.96 %
EW precision constraints	9.73 %	9.81 %	9.68 %	10.51 %
All non-DM constraints	32.3 %	35.1 %	31.8 %	41.10 %
Overall ATLAS Run 2 + non-DM external	40.0 %	38.7 %	65.2 %	55.8 %
Relic density constraint	96.03 %	12.81 %	0.25 %	0 %
Direct detection constraints	83.80 %	67.72 %	24.37 %	64.71 %
All external constraints	98.7 %	79.8 %	48.2 %	78.73%
Overall ATLAS Run 2 + external	98.81 %	80.90 %	81.03 %	82.50 %

Table 4.13.: Percentage of models excluded by each result, for each LSP type and scan.

Bino LSP

Figure 4.48 shows the mass-spectrum for four benchmark models with a bino-like LSP (from the Bino-DM scan) which pass all constraints and are not excluded by ATLAS Run

2 analyses. The $\tilde{\chi}_1^0$, $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ masses of all of these models lie within the published ATLAS simplified model contours for at least one of the SUSY searches considered in this analysis.

Figure 4.48a shows a model in the Z/h funnel region. This model differs from typical simplified models as the $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ higgsino fraction is greater than 98%. Additionally, this scenario has $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h) \approx \text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z) \approx 50\%$. Therefore scenarios such as this will have smaller signal yields than typical simplified models due to (i) the higgsino-like $\tilde{\chi}_1^\pm / \tilde{\chi}_2^0$ production cross-sections being smaller than wino-like, (ii) the branching ratios to the typical $\tilde{\chi}_2^0$ decay modes being half the usual simplified model assumption of 100%. The best analysis for this model is 1Lbb, with a final $\text{CL}_s^{\text{Obs.}}$ value of 0.33. The 1Lbb analysis targeted the $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h$ decay, so perhaps a statistical combination with 2L2J, which targeted the $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z$ decay, would provide sensitivity to such a model.

Figure 4.48b shows a model with a wino $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^0$ with a large mass splitting to the bino LSP. This scenario has a light A and H such that the dark matter relic density is reduced via the A/H funnel self-annihilation mechanisms, $\tilde{\chi}_1^0 \tilde{\chi}_1^0 \rightarrow A/H$. Similar to the Z/h funnel model, this model has mixed $\tilde{\chi}_2^0$ decay modes, with $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h) = 75\%$ and $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z) = 25\%$. The model lies just beyond the 1Lbb region of sensitivity, so despite the $\tilde{\chi}_2^0$ decay to h being dominant, the 2L2J search (which targeted the $\tilde{\chi}_2^0$ decay via Z) has the best sensitivity. As before, a statistical combination may improve the sensitivity here.

Finally, two scenarios with small mass splittings between the $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ are shown in Figures 4.48c and 4.48d. As discussed in Section 4.4.1, these scenarios have a smaller branching fraction for the $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \ell^+ \ell^-$ three-body decay than the simplified models used in the original Compressed analysis interpretation. They are instead dominated by the radiative decay $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \gamma$. Therefore these scenarios are not excluded despite having masses within the Compressed analysis exclusion contour. In the case of Figure 4.48c, the $\tilde{\chi}_2^0$, $\tilde{\chi}_2^\pm$ and $\tilde{\chi}_4^0$ are light enough to be within reach of searches such as FullHad and 2L2J. Indeed, 2L2J provides the best sensitivity to this scenario, with a final $\text{CL}_s^{\text{Obs.}}$ value of 0.2. Perhaps a dedicated search optimised to scenarios like this could improve on this sensitivity.




Figure 4.48.: Mass-spectrum for four benchmark models with a bino-like LSP which pass all constraints and are not excluded despite having a mass-spectrum within published ATLAS simplified model contours. Produced using PySLHA [167].




Figure 4.49.: Mass-spectrum for two benchmark models with a higgsino-like LSP which pass all constraints and are not excluded despite having a mass-spectrum within published ATLAS simplified model contours. Produced using PySLHA [167].

Higgsino LSP

Figure 4.49 shows the mass-spectrum of two benchmark models with a higgsino-like LSP (from the EWKino scan) which pass all constraints and are not excluded by ATLAS Run 2 analyses. The $\tilde{\chi}_1^0$ mass and $\tilde{\chi}_2^0$ mass splitting of these models is within the published simplified model contours for the ATLAS 3L (offshell) and Compressed searches.

As with the compressed bino LSP model discussed previously, the $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_2^0 \ell^+ \ell^-$ decay mode has a smaller branching fraction for these models than the Compressed analysis simplified models, and the radiative decay is the dominant decay mode. The $\tilde{\chi}_2^0$ decay to $\tilde{\chi}_1^\pm$ via an offshell W boson also has a non-zero branching ratio for these scenarios. Otherwise, the higgsino/bino/wino are well separated in these models, making them fairly simplified-model-like. The $\tilde{\chi}_3^0$ and $\tilde{\chi}_2^\pm$ have masses under a TeV and large mass-splittings with the other electroweakinos for both models. Therefore sensitivity could be achieved to the production of these heavier electroweakinos and their high p_T decay products.

4.8.3. Outlook

The results presented in this chapter provide a legacy of the Run 2 ATLAS searches for electroweak production of supersymmetric particles. The suite of ATLAS searches was originally optimized for and interpreted in simplified model scenarios, focusing on a particular production and decay mode. The interpretation of these searches in the

pMSSM shown here vindicates this simplified model approach. Collectively, the searches have sensitivity to a broad range of pMSSM scenarios, including those that are not excluded by measurements of flavour and EW precision observables or dark matter direct detection experiments. A significant improvement is seen compared to the previous ATLAS pMSSM interpretations performed with Run 1 search results in References [128] and [129]. In Reference [129], no sensitivity was seen to bino-LSP scenarios outside the funnel regions. These updated results show the funnel regions almost entirely excluded and new sensitivity to electroweak scenarios with compressed mass-spectra.

However, many scenarios remain uncovered – even those with masses well within published limits. Higgsino-LSP scenarios with $\mathcal{O}(\text{GeV})$ mass splittings, for example, are particularly under-covered. It is important to target these scenarios with future searches. Run 3 of the LHC will approximately double the size of the ATLAS dataset, and re-running similar analyses on the new data will provide some modest improvement to the sensitivity seen previously. But, it may be time to move beyond the simplified model paradigm and perform targeted searches for the most likely remaining scenarios – those that satisfy all existing constraints and have evaded ATLAS sensitivity. These may have non-simplified spectra and mixed decay modes. The pMSSM can provide a framework of signal models within which to design these searches. Indeed, perhaps the remaining viable models in the scans presented here can be used as benchmarks by which to interpret, or even optimise, new searches. This could avoid over-tuning to non-viable simplified models, and maximise sensitivity to the most likely scenarios.

Finally, there are several avenues to develop this study of ATLAS sensitivity to pMSSM scenarios. Direct detection experiments place strong constraints on bino-LSP scenarios in the Z/h funnel regions (even those with a satisfactory dark matter relic density). This study could be improved by providing a larger sample of bino-LSP models satisfying this constraint such that a more granular analysis of the ATLAS sensitivity to surviving scenarios could be performed.

Also, the interpretation of models with low-mass and compressed $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$ and $\tilde{\chi}_1^0$ is limited by several simplifications outlined in Section 4.2.3, such as low event-statistics and simplifications in the evaluation of the Compressed analysis. A future iteration of this scan should strive to improve this, as compressed scenarios are the least well covered making them the most important to understand and target with future searches.

The scan and interpretation here focused specifically on the production of electroweakinos. Of course, if realised in nature, supersymmetry may also provide sleptons, squarks

and gluinos at low mass. Indeed a small mass splitting between a bino-LSP and slepton or squark provides another co-annihilation mechanism through which the dark matter relic density can be reduced. The inclusion of low-mass sleptons can also produce a value of $g - 2$ compatible with the current experimental measurement [22, 23]. Scans of the pMSSM which include strongly produced sparticles, and sleptons, are under development and will include reinterpretation of the full suite of ATLAS Run 2 searches for both electroweak and strong production of supersymmetric particles.

The remaining chapters of this thesis consider two approaches which could be used to enhance the sensitivity of ATLAS to the remaining viable pMSSM scenarios. Chapter 5 will demonstrate the potential for statistical combinations to significantly extend the range of exclusion in simplified model scenarios. Statistical combinations such as this are not performed when evaluating the pMSSM exclusion in this chapter. Developing the workflow to perform such statistical combinations may provide sensitivity to some of the missed scenarios with mixed branching ratios that no individual analysis can exclude alone. Finally, Chapter 6 will consider the potential sensitivity gain from improved p_T^{miss} reconstruction performance.

Chapter 5.

Statistical combination of searches for chargino production

5.1. Introduction

In Chapter 4, a suite of ATLAS Run 2 searches for supersymmetry (SUSY) was introduced which targets a range of simplified SUSY scenarios. Some of these searches targeted the same simplified models but different final states, so have both unique and overlapping sensitivity in different regions of the same simplified model parameter space. A statistical combination of such searches can extend and deepen the existing sensitivity offered by each analysis alone.

This chapter presents the statistical combination of searches which targeted the direct production of a pair of charginos, followed by the decay of each chargino to the lightest neutralino via emission of a W boson: $pp \rightarrow \tilde{\chi}_1^+ \tilde{\chi}_1^- \rightarrow W^+ \tilde{\chi}_1^0 W^- \tilde{\chi}_1^0$. This process is shown in Figure 5.1. The interpretation is for a simplified model where the $\tilde{\chi}_1^\pm$ are purely wino, $\tilde{\chi}_1^0$ is purely bino and the branching ratio of the $\tilde{\chi}_1^\pm \rightarrow W^\pm \tilde{\chi}_1^0$ decay is 100%. The only free parameters are the masses of $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_1^0$, so signal points will be referred to using the notation (X, Y) where $X = m(\tilde{\chi}_1^\pm)$ and $Y = m(\tilde{\chi}_1^0)$ with units in GeV.

These results form part of the ATLAS Run 2 Electroweak SUSY Combinations effort, which also includes combinations of analyses targeting wino-like $\tilde{\chi}_1^+ \tilde{\chi}_2^0$ production followed by WZ and Wh decays, Higgsino GGM models, and slepton pair production. These results update the previous combination of Run 1 ATLAS electroweak SUSY analyses [183], providing the strongest ATLAS exclusion limits on the dominant electroweak SUSY production mechanisms.




Figure 5.1.: Feynman diagram for the direct production of a pair of charginos, followed by decays to the lightest neutralino via emission of a W boson.

5.2. Searches for $\tilde{\chi}_1^+ \tilde{\chi}_1^-$ production

The targeted simplified model produces a pair of lightest neutralinos plus a pair of W bosons which can subsequently decay either leptonically ($W \rightarrow \ell \bar{\nu}_\ell$) or hadronically ($W \rightarrow qq'$). Hence, depending on the decay mode, the final state consists of missing transverse momentum (p_T^{miss}) from the neutralinos, plus either zero leptons and hadronic jets, one lepton and hadronic jets, or 2 leptons and no hadronic jets. Each analysis to be combined targets one of these three cases:

- The first [60] and second [123] wave ‘two leptons plus zero jets’ (2L0J) analyses target the fully leptonic case, $WW \rightarrow \ell \bar{\nu}_\ell \ell' \bar{\nu}_{\ell'}$.
- The one lepton (1L) analysis [184] targets the semi-leptonic case $WW \rightarrow \ell \bar{\nu}_\ell qq'$.
- The fully hadronic (FullHad) analysis [139] targets the fully hadronic case $WW \rightarrow q_1 q'_1 q_2 q'_2$.

Note that 2L0J (wave 1) and FullHad were seen previously in Chapter 4 where they were interpreted in pMSSM scenarios. The 1L and 2L0J (wave 2) analyses on the other hand were not available in time for the study in Chapter 4. Figure 5.2 shows the separate




Figure 5.2.: The original expected (dashed lines) and observed (full lines) exclusion contours for each of the analyses used in this combination.

95% CL exclusion limits for each of the analyses in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane. There is some overlap between the analyses as well as unique sensitivity.

FullHad can access the largest $\tilde{\chi}_1^\pm$ masses as the hadronic decay of the W boson has a larger branching ratio than the leptonic decay mode (68% vs. 32%) which compensates for the smaller production cross-sections at higher mass. Since FullHad requires no final state leptons, the p_T^{miss} trigger is used along with an offline selection $p_T^{\text{miss}} > 200$ GeV such that all events have approximately 100% trigger efficiency. Additionally, to discriminate from large hadronic Standard Model backgrounds, boosted large radius (large- R) jets with $p_T > 200$ GeV are required. These high p_T selections on the final-state objects preclude sensitivity to small or moderate mass splitting between the $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_1^0$ where the phase-space is more restricted.

The 2L0J wave 1 analysis on the other hand cannot access larger $\tilde{\chi}_1^\pm$ masses (where the production cross-section is smaller) due to the smaller branching ratio of the leptonic W decay mode. Events are selected with the di-lepton trigger and require two reconstructed leptons with $p_T > 10$ GeV and $p_T^{\text{miss}} > 100$ GeV. These thresholds are much lower than the FullHad analysis, therefore lower mass signal points with smaller mass splitting are accessible. The 2L0J wave 2 analysis uses a boosted decision tree to target even lower mass and moderately compressed scenarios.

The 1L analysis targets one hadronic and one leptonic W decay, providing sensitivity between 2L0J and FullHad on the $m(\tilde{\chi}_1^\pm)$ axis. Events are selected with a single lepton trigger and the requirement of one reconstructed lepton with $p_T > 25$ GeV. One to three jets with $p_T > 30$ GeV and one large-R W -tagged jet with $p_T > 250$ GeV are also required, along with $p_T^{\text{miss}} > 200$ GeV. Lower p_T thresholds can be applied to leptons than jets, therefore the use of a final-state lepton in the 1L analysis allows access to points with a smaller mass splitting compared to FullHad. However, the lepton p_T requirement is higher than that in the 2L0J analysis. This, along with the requirement of a very high p_T jet and large p_T^{miss} , reduces the sensitivity to the lower mass and moderately compressed scenarios (which have a more restricted phase-space) that 2L0J can access.

The statistical combination aims to fill the gaps in sensitivity between these analyses and extend the reach to smaller mass splittings.

5.3. Method

5.3.1. Performing a statistical combination

A combined statistical fit is performed using the statistical analysis package PyHF [114, 115] v0.7.0 with the Jax [185] backend and SciPy [107] optimizer. Other backends and optimizers were also tested and confirmed to produce almost identical results. Jax and SciPy produced the most stable combination with the fewest failed fits.

The key ingredient for each signal point and analysis is a JSON-serialized workspace, introduced in Chapter 3. The workspaces for each signal point and analysis contain the background and signal model estimation from each individual analysis. Preserved workspaces from the original analyses are used without modification (unless specified) where available. Some new signal points were generated to improve the harmonization of the grid points across the analyses and increase the density of points in the high-mass region. New signal point workspaces were generated with the same set-up described in each analysis publication, and processed using RECAST [118].

The final outputs of the combination are expected and observed CL_s values and upper limits on the signal strength (μ_{SIG}). A point is considered excluded if $\text{CL}_s < 0.05$.




Figure 5.3.: Signal point workspaces available for each analysis.

5.3.2. Analyses combined for each signal point

As the individual analyses target different regions of the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ plane they do not have a perfectly harmonized set of signal points, though as mentioned in Section 5.3 some new signal points were generated to improve the harmonization at high mass. The points available for each of the analyses are shown in Figure 5.3a.

However, not all of these points are used when calculating the final combined exclusion contour. Discontinuities in the exclusion of neighbouring points can arise when a point for which just one analysis workspace is available neighbours a point for which multiple analysis workspaces are available, as the latter point receives an improvement in CL_s from the combination. This causes issues in the calculation of the exclusion contour, so some points which have been processed for just one analysis are removed from the final results. The points which are included are shown in Figure 5.3b. The majority of points that are not included in the combination lie on the high $\tilde{\chi}_1^\pm$ mass edge of the 2L0J wave 1 contour (around $m(\tilde{\chi}_1^\pm) = 400$ GeV) as a workspace is only available for the 2L0J wave 1 analysis. Additionally, several other points are not included in the combination where a workspace is only available for either 1L or FullHad.

The analyses entering the combination must be statistically independent and not overlap in their event selection. Orthogonality has been built into the electroweak SUSY analyses by using common baseline lepton definitions and requiring orthogonality in lepton multiplicity for each analysis. The baseline lepton definitions used in each electroweak SUSY analysis are:

- Electrons: $p_T > 4.5$ GeV, $|\eta| < 2.47$, LooseAndBLayerLLHID [186], $|z_0 \sin \theta| < 0.5$ mm.
- Muons: $p_T > 3$ GeV, $|\eta| < 2.7$, Medium ID [187], $|z_0 \sin \theta| < 0.5$ mm.

Hence the signal regions for FullHad (no leptons), 1L (one lepton) and 2L0J (2 leptons) are already guaranteed to be orthogonal before any complex analysis selections are applied. Note that two iterations of the 2L0J analysis are considered here (wave 1 and wave 2), both of which require 2 leptons and have overlapping selections. For any given point only one of them can be used in the combination. For each point, the 2L0J analysis workspace which results in the lowest expected CL_s when fit on its own is used in the combination. After this selection procedure, there are no points that overlap with 1L or FullHad for which the 2L0J wave 2 workspace is used in the combination.

The overlap between control regions for the analyses to be combined is also tested by comparing the events selected in data. Only two pairs of analyses are found to have overlapping data events:

1. 2L0J wave 1 and wave 2 – only the analysis with the best expected CL_s is used as discussed above.
2. 9 events overlap between 2L0J wave 2 and 1L. However, the overlap is for control regions that are not used for the simplified model considered for this combination. Furthermore, as discussed above 2L0J wave 2 is not combined with 1L for any signal points since they target very different signal mass scenarios.

Note that workspaces are available for the (400, 0), (500, 0) and (600, 0) points for FullHad and 1L while the 2L0J analysis instead has workspaces for the (400, 1), (500, 1) and (600, 1) points. These 2L0J workspaces are treated as (400, 0), (500, 0) and (600, 0) points for the purposes of the combination so that a full combination of all three analyses can be produced for these points.

5.3.3. Correlating systematic uncertainties



Some systematic uncertainties are shared between analyses, and should therefore be correlated in the combined statistical fit. “Correlating” systematic uncertainties from different analyses means that they correspond to the same nuisance parameter in the statistical fit. However, where there are differences in the configuration of object definitions the correlation of associated uncertainties becomes non-trivial. A conservative

approach is taken here and uncertainties are only correlated where they have identical configurations and are associated with identical object definitions across analyses.

The lepton, pile-up re-weighting, p_T^{miss} , luminosity and signal cross-section uncertainties are correlated across all of the analyses. The Jet Energy Scale (JES) uncertainties for small-radius (small-R) jets are correlated across 2L0J and 1L as both analyses used the same configuration. The FullHad analysis only used small-R jets for the p_T^{miss} and $\min\Delta\phi(j, p_T^{\text{miss}})$ calculation and applied a simplified JES uncertainty configuration which precludes correlation with the 2L0J and 1L small-R JES uncertainties. The JES systematic uncertainties associated with large-R jets, and all Jet Energy Resolution (JER) uncertainties, are not correlated across any of the analyses due to differences in the configurations used. This is unfortunate given the overlap in sensitivity between 1L and FullHad and the importance of jets in both analyses.

Some theoretical uncertainties are also shared between analyses. For FullHad and 1L, the following theoretical uncertainties are correlated: the factorisation, renormalisation, PDF, matching (CKKW) and resummation uncertainties of the Z+jets and W+jets backgrounds. For 2L0J and 1L, the diboson renormalisation, factorisation and PDF uncertainties are correlated. For some background processes different (versions of) Monte Carlo generators were used for different analyses, and the corresponding uncertainties are left uncorrelated for the combination.

Figure 5.4 shows the absolute and relative difference between the combination CL_s values with and without the experimental systematics correlated. Figure 5.5 shows the absolute and relative difference between the combination CL_s values with and without the theoretical systematics correlated (but correlating experimental systematics in both cases). For the majority of points, the impact of correlating the systematics is negligible, though a handful of points around (500, 200) do see a 10-20% reduction in CL_s values when correlating the systematics. The final results will be presented for a combination with both uncorrelated and fully correlated (both experimental and theoretical) systematics.



5.4. Validation

5.4.1. Validation of analysis inputs

The original analyses used HistFitter [188] for the computation of their CL_s values, while the combination is performed using PyHF as outlined in Section 5.3. It has been verified that the difference in the CL_s values obtained using PyHF instead of HistFitter is negligible. This can be seen in Figures 5.6 and 5.7, which show the original CL_s values for each analysis, and the difference for the value obtained by re-running the fit with PyHF. No differences within two decimal places are seen for FullHad. There are differences for 2L0J wave 1 and 2 for several points, though these are mostly outside each analysis' area of sensitivity, and do not impact the exclusion contours. The (700, 1) point has a particularly large discrepancy for 2L0J wave 1, which may be due to an incorrectly preserved workspace. However the nearby point (700, 0) is already excluded by FullHad, therefore the (700, 1) point is not needed for the combination and is not included when producing the final results.

For 1L, a large discrepancy is seen for the (500, 300) point. However this point is not excluded with a large CL_s value in both cases, so this does not impact the exclusion contour. Note that recalculated CL_s values for the (300, 0) point for 1L are missing because the fit failed. This point is removed and does not affect the final combination exclusion contour as the nearby point (300, 1) is excluded by 2L0J.

The workspaces for the (700, 200) and (500, 300) points for the 1L analysis contained deviations of greater than 1000% for some signal systematic uncertainties in a control region. This is unphysically large, and for (700, 200) leads to the unexpected exclusion of this point when running the combined fit. Removing these unphysical deviations (by setting them to 100%) solves this problem, and produces a more sensible CL_s result. The reported results for these points use these modified workspaces.

5.4.2. Fit stability

The stability of the combination fits must be checked to validate that the combination is not introducing additional pulls or constraints in the nuisance parameters. To access the uncertainties on the parameters required to create pull plots, the combination is re-run using the Minuit [106] optimiser (with Jax backend). It has been verified that the CL_s




Figure 5.6.: The original CL_s values (produced with HistFitter) for 2L0J wave 1 and 2, and the difference for the recalculated values (produced with PyHF).




Figure 5.7.: The original CL_s values (produced with HistFitter) for the 1L and FullHad analyses, and the difference for the recalculated values (produced with PyHF).

values produced from combinations using the Minuit optimiser match those produced using the SciPy optimiser.

To assess the impact of the combination on pulls of nuisance parameters, pull plots are made for three points which combine different sets of analyses. Figures 5.8, 5.9 and 5.10 show pull plots for the expected exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (300, 100)$, $(600, 200)$ and $(900, 300)$ points respectively. Each of these points combined a different set of analyses, as indicated in the captions. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are displayed, with the original analysis pull in black and the combination pull in red. For completeness, the pull plots for the observed fit are included in Appendix A, along with the systematic uncertainty correlation matrix.

No significant pulls or constraints are observed and there are no large changes in the combination pulls relative to each individual analysis. It is concluded from this that the combination is not introducing any additional pulls or constraints. The pull plots for several other mass points have also been checked, and the same conclusions drawn.

5.5. Results

The exclusion contours are shown in Figures 5.11 and 5.12 for combinations performed with both uncorrelated and correlated experimental and theoretical systematic uncertainties. The combination has filled the exclusion gap between the 1L and FullHad observed/expected exclusion contours and reduced the large expected uncertainty bands for those analyses. The exclusion has improved significantly above the FullHad exclusion contour, where neither 1L nor FullHad could reach alone. The improvement is more moderate in the region between the 2L0J and 1L contours around $300 \text{ GeV} < m(\tilde{\chi}_1^\pm) < 500 \text{ GeV}$, but larger (and more compressed) $\tilde{\chi}_1^0$ masses are excluded than by either analysis alone. Correlating systematic uncertainties between the analyses slightly improves the overall sensitivity, particularly in the central region $m(\tilde{\chi}_1^\pm) \approx 500 \text{ GeV}$, but overall does not have a large impact.

A kink in the $-1\sigma_{\text{exp.}}$ uncertainty band can be seen around $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (500, 100)$ – in between the regions covered by 1L and FullHad alone. The combination has improved the sensitivity within the large 1L and FullHad $-1\sigma_{\text{exp.}}$ uncertainty bands, and this kink is in the middle where the combination is weakest. The discontinuous shape is likely




Figure 5.8.: Pull plot from an expected exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (300, 100)$ mass point, which combines the 2L0J and 1L analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pulls are shown in red, the original analysis pulls in black.




Figure 5.9.: Pull plot from an expected exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (600, 200)$ mass point, which combines the 2L0J, 1L and FullHad analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pulls are shown in red, the original analysis pulls in black.




Figure 5.10.: Pull plot from an expected exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (900, 300)$ mass point, which combines the 1L and FullHad analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pulls are shown in red, the original analysis pulls in black.




Figure 5.11.: Combined expected exclusion limit overlaid with the exclusion contours from each individual analysis, for a combination performed (a) without correlating systematics uncertainties and (b) with experimental and theoretical systematics correlated.




Figure 5.12.: Combined expected exclusion limit overlaid with the expected CL_s values for each signal point (note for clarity CL_s values are not displayed for all available points). (a) shows the results without correlating systematics uncertainties and (b) with experimental and theoretical systematics correlated.

an artefact due to a lack of point density nearby. In particular, the points (600, 50) and (700, 50) are unavailable for the combination (as a workspace was only available for FullHad for these points), which may be impacting the interpolation for this uncertainty band.




Figure 5.13.: The upper limits on the signal strength μ_{SIG} set by the combination (with correlated systematics), and the relative change with respect to using the analysis with the best expected limit at each point with no combination. Note that for clarity the results are not displayed for all available points.

The depth of exclusion is also checked by studying the limits on the signal strength, μ_{SIG} , as seen in Figure 5.13. Figure 5.13b shows the percentage reduction in the signal strength upper limit for the combination relative to the analysis with the best expected limit. An improvement of 10 – 30% can be seen in the region of sensitivity.

5.6. Conclusions

A statistical combination of three searches (1L, 2L0J, FullHad) for charginos has been performed. The results are interpreted as exclusion limits on a simplified model of pure-wino chargino pair production, with each chargino decaying to a pure-bino lightest neutralino via the emission of a W boson.

The combination significantly improves the exclusion in the $\tilde{\chi}_1^\pm$ - $\tilde{\chi}_1^0$ mass plane and the upper limits on the signal strength. This combination is only possible due to the orthogonal event selections of these three analyses. This orthogonality has been built into the ATLAS electroweak supersymmetry searches by design with a view to this and other combinations being performed. This choice is vindicated by the improvements shown here

and motivates a similar approach being taken for Run 3 searches. Closer harmonization of signal point grids across different searches would also aid future statistical combinations.

Finally, it is emphasized that the exclusion reported in this chapter is for simplified models only. As demonstrated in Chapter 4, it is still possible for charginos and neutralinos with masses well within the reported contour to have evaded detection if supersymmetry, as realised in nature, does not satisfy the assumptions of the simplified model considered. Chapter 4 included a pMSSM interpretation of the 2L0J and FullHad searches discussed in this chapter. However, the 1L search, and statistical combinations of the searches, were not included. The studies presented here validate the procedure for a combination of 1L, 2L0J and FullHad. This opens the door to a future iteration of the pMSSM scan which includes combinations of these three analyses, along with combinations of analyses targeting other electroweak production mechanisms, such as $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$ pair production with decays via WZ or Wh and slepton production. This could improve the sensitivity to some of the pMSSM models on the edge of exclusion in Chapter 4 and would provide the most complete picture of the status of electroweak supersymmetry scenarios following the Run 2 ATLAS searches.

Chapter 6.

METNet: A combined p_T^{miss} working point using a neural network

6.1. Introduction

Missing transverse momentum (p_T^{miss}) [94], introduced in Section 3.4.6, is a key discriminating observable in many ATLAS searches for supersymmetry. It represents the total transverse momentum of undetected particles produced in a proton–proton collision. Dark matter candidates (such as the lightest neutralino in the pMSSM) would escape ATLAS undetected if produced in a proton–proton collision, resulting in a non-zero value of p_T^{miss} . Standard Model neutrinos are also invisible to ATLAS and result in p_T^{miss} . However, finite detector resolution, finite detector acceptance, and contamination of particles produced in pile-up interactions can all result in the reconstructed p_T^{miss} differing from the true value of p_T^{miss} ($p_T^{\text{miss, True}}$). A non-zero reconstructed p_T^{miss} for an event with zero $p_T^{\text{miss, True}}$ is referred to as *fake* p_T^{miss} .

The jet contribution to the p_T^{miss} calculation is particularly sensitive to pile-up contamination and, as discussed in Section 3.4.6, ATLAS employs a suite of jet selection working points for p_T^{miss} reconstruction. The optimal working point (that which gives a reconstructed p_T^{miss} closest to $p_T^{\text{miss, True}}$) for a given event is dependent on the pile-up conditions and topology of the event. This is demonstrated in Figure 6.1 which shows the root-mean-square of the difference between the reconstructed value of $p_{x,y}^{\text{miss}}$ and $p_{x,y}^{\text{miss, True}}$ (referred to as root-mean-squared error or RMSE) in bins of average number

of interactions per bunch crossing $\langle\mu\rangle$ for the *Tight*, *Loose*, *Tenacious* and *TightClus* working points. Here *TightClus* uses a cluster-based soft term (but the same jet selections as *Tight*) while the other working points use a track-based soft term. For $t\bar{t}$ events (left), the *TightClus* and *Loose* working points have the smallest RMSE at low $\langle\mu\rangle$. But as $\langle\mu\rangle$ increases, corresponding to increased in-time pile-up, the performance of the *TightClus* and *Loose* working points deteriorate in favour of *Tenacious* (which has the strictest jet requirements). On the other hand for $Z \rightarrow \mu\mu$, a topology which produces fewer hard-scatter jets than $t\bar{t}$, the *Tenacious* working point is the optimal choice for almost the entire range of $\langle\mu\rangle$.




Figure 6.1.: Root-mean-square of the difference between reconstructed and true p_T^{miss} in bins of average number of interactions per bunch crossing, for (a) $t\bar{t}$ and (b) $Z \rightarrow \mu\mu$ events. The lower panel shows the ratio with respect to the *Tight* working point. A hatched band indicates the statistical uncertainty for the *Tight* working point (which is negligible in most bins).

Currently, an ATLAS analysis will choose one of these working points (most analyses use *Tight*) and use it to reconstruct the p_T^{miss} in every event, which is sub-optimal. This motivates an event-by-event combination of the working points to produce a single p_T^{miss} definition that is suitable for all events. In this chapter, two neural network (NN) architectures which produce such a combination are presented, and their performance is compared to each other and the existing p_T^{miss} working points. Given the historical convention of calling p_T^{miss} ‘MET’, the neural network approaches are collectively named ‘METNet’. The two architectures will be distinguished using the names METNet_R and METNet_C:

1. METNet_R: A regression targeting $p_x^{\text{miss, True}}$ and $p_y^{\text{miss, True}}$.
2. METNet_C: A classifier targeting the optimal working point for each event.

In each case, the neural network combines complementary information from each working point with additional information characterising event topology to produce a p_T^{miss} prediction with improved resolution and pile-up resilience compared to current working points.

This chapter presents the machine-learning methods used to define the neural networks and MC-based studies of their performance compared to current ATLAS p_T^{miss} working points. The chapter is structured as follows: In Section 6.2, the simulated samples used to train and test the NNs are described along with the object definitions and event selections applied. The input features, architectures and data pre-processing steps used for METNet_R and METNet_C are described in Section 6.3. Section 6.4 presents the performance of the p_T^{miss} prediction of each NN compared to current ATLAS p_T^{miss} working points, before discussing which of the two architectures is the best choice for use in future analyses. In particular, the improvements to signal vs. background separation for pMSSM signal points from Chapter 4 is discussed. Finally, conclusions and suggestions for further study are outlined in Section 6.6.

6.2. Training and testing data

6.2.1. Simulated event samples

The neural networks are trained and tested using simulated event samples for a variety of production processes. Section 3.5.1 introduced the full and fast simulation procedures for simulating the impact of the ATLAS detector. The Standard Model processes used in this study are passed through a full simulation of the ATLAS detector. The effect of pile-up in the same and neighbouring bunch crossings are modelled by overlaying the simulated hard-scatter event with inelastic proton–proton events generated with PYTHIA8.186 [189] using the NNPDF2.3 set of parton distribution functions (PDF) [190] and the A3 set of tuned parameters [191]. The MC samples are reweighted so that the distribution of the average number of interactions per bunch crossing reproduces the observed distribution in data.

Table 6.1.: Simulated event samples with the corresponding matrix element (ME) and parton shower (PS) generators, cross-section order in α_s used to normalise the event yield, underlying-event tune and the generator PDF sets used.

Physics process	Generator (ME)	Parton shower	Cross-section	Tune	PDF (ME)	PDF (PS)
$t\bar{t}$	POWHEG Boxv2 [195–197]	PYTHIA8.230 [99]	NNLO+NNLL [198]	A14	NNPDF3.0NLO [199]	NNPDF2.3LO
VV	POWHEG Boxv2	PYTHIA8.210	NLO [193]	AZNLO [200]	CT10 NLO [201]	CTEQ6L1 [202]
VBF $V \rightarrow XX$	SHERPA2.2.1 [100, 203–205]	SHERPA2.2.1	NNLO [206]	SHERPA default	NNPDF3.0NNLO	NNPDF3.0NNLO
$Z(\rightarrow \ell\ell)+\text{jets}$	SHERPA2.2.1	SHERPA2.2.1	NNLO	SHERPA default	NNPDF3.0NNLO	NNPDF3.0NNLO
pMSSM	MADGRAPH 2.9.5 [164]	PYTHIA8.306	NLO	A14	NNPDF2.3LO	NNPDF2.3LO

Additionally, a selection of pMSSM model points from the scans in Chapter 4 are used to test the NNs. The corresponding event samples are processed with a fast simulation of the ATLAS detector. Full details of how these samples are produced can be found in Section 4.2. Note that the pMSSM signal includes electroweak production only and is generated to leading order (LO) in QCD, therefore LO PDF sets are used.

The samples used are listed in Table 6.1 along with the relevant parton distribution function (PDF) sets, the configuration of underlying-event and hadronisation parameters (tune), and the cross-section order in α_s used to normalise the event yields for these samples. Further information on the ATLAS simulations of $t\bar{t}$, multiboson and boson plus jet processes can be found in the relevant public notes [192–194].

6.2.2. Object definitions

The objects in the simulated samples are reconstructed through the procedure outlined in Section 3.4. Table 6.2 summarises the specific kinematic selection and quality criteria used for the electrons, muons and jets. Both baseline and signal objects are defined. Baseline objects are used in the p_T^{miss} calculation while signal objects must satisfy more stringent selection criteria and are used for the final topology selections. Note that hadronic τ -lepton decay products are treated as jets.

The definitions of the p_T^{miss} working points used in this study are listed in Table 6.3. Note that this includes several additional working points to those introduced in Section 3.4.6. All working points use a track-based soft term other than *TightClus* which uses a cluster-based soft term. This includes neutral objects but at the expense of increased pile-up contamination. The *Track* working point is also considered which builds

Table 6.2.: Summary of object definitions. Baseline objects are used in the p_T^{miss} calculation while signal objects are used in the event selections.

Property	Object definitions	
	Baseline	Signal
Electrons		
Kinematic	$p_T > 10 \text{ GeV}, \eta < 2.47$	$p_T > 25 \text{ GeV}$
Identification	<i>Loose</i> likelihood [207]	<i>Tight</i> likelihood
Isolation	Loose if $p_T < 200 \text{ GeV}$ else <i>HighPtCaloOnly</i> [207]	Loose if $p_T < 200 \text{ GeV}$ else <i>HighPtCaloOnly</i> [207]
Impact parameter	$ z_0 \sin \theta < 0.5 \text{ mm}$	$ z_0 \sin \theta < 0.5 \text{ mm}, d_0 /\sigma(d_0) < 5$
Muons		
Kinematic	$p_T > 10 \text{ GeV}, \eta < 2.7$	$p_T > 25 \text{ GeV}, \eta < 2.7$
Identification	<i>Medium</i> [87]	<i>Medium</i>
Isolation	Loose [208]	Loose
Impact parameter	$ z_0 \sin \theta < 0.5 \text{ mm}$	$ z_0 \sin \theta < 0.5 \text{ mm}, d_0 /\sigma(d_0) < 3$
Jets		
Kinematic	$p_T > 20 \text{ GeV}, \eta < 4.5$	$p_T > 20 \text{ GeV}, \eta < 2.8$
Clustering	Anti- k_t R=0.4 PFlow	Anti- k_t R=0.4 PFlow
Pile-up mitigation	–	JVT > 0.5 for $p_T < 60 \text{ GeV}, \eta < 2.4$ fJVT < 0.4 for $ \eta > 2.5, p_T < 120 \text{ GeV}$
b -tagging	–	$p_T > 20 \text{ GeV}, \eta < 2.5$, DL1 77% eff WP [209]

p_T^{miss} entirely from tracks (rather than a sum of hard and soft objects). This provides very good pile-up rejection but ignores all neutral objects produced in the hard scatter. Note that the electron and muon p_T^{miss} terms are identical for each working point (other than *Track*), as only the jet selections differ.

6.2.3. Training and testing datasets

A variety of topologies are included in the training dataset such that the NNs can generalise well. The training set includes a mixture of $t\bar{t}$, $ZZ \rightarrow \nu\nu\ell\ell$ and $WZ \rightarrow \ell\nu\ell\ell$ events — topologies which produce real p_T^{miss} . Topologies with no real p_T^{miss} , such as $Z \rightarrow \ell\ell$, are excluded from the training set, as they introduce a severe bias towards zero in METNet_R's p_T^{miss} prediction. Initial studies were performed with a training dataset

Table 6.3.: Jet selections for p_T^{miss} reconstruction working points used.

Working point	Selections			fJVT for jets with
	p_{T} [GeV] for jets with: $ \eta < 2.4$	$2.4 < \eta < 4.5$	JVT for jets with $ \eta < 2.4$	$2.5 < \eta < 4.5$ and $p_{\text{T}} < 120$ GeV
<i>Loose</i>	> 20	> 20	> 0.5 for $p_{\text{T}} < 60$ GeV jets	–
<i>Tight</i>	> 20	> 30	> 0.5 for $p_{\text{T}} < 60$ GeV jets	–
<i>TightClus</i>	> 20	> 30	–	–
<i>TightfJVT</i>	> 20	> 30	> 0.5 for $p_{\text{T}} < 60$ GeV jets	< 0.4
<i>Tighter</i>	> 20	> 35	> 0.5 for $p_{\text{T}} < 60$ GeV jets	–
<i>Tenacious</i>	> 20	> 35	> 0.91 for $20 < p_{\text{T}} < 40$ GeV jets	< 0.5
			> 0.59 for $40 < p_{\text{T}} < 60$ GeV jets	
			> 0.11 for $60 < p_{\text{T}} < 120$ GeV jets	
<i>Track</i>	N/A – Built from tracks rather than a sum of hard and soft objects.			

consisting solely of $t\bar{t}$ events, but the inclusion of di-boson topologies improved the performance of the NN’s on unseen test topologies.

Each NN is trained on 8 million events of which about 40% are $t\bar{t}$, 30% $ZZ \rightarrow \nu\nu\ell\ell$ and 30% $WZ \rightarrow \ell\nu\ell\ell$. These events do not have MC generation event weights applied, nor are they weighted to account for their cross-sections. An additional 800,000 events with the same topology proportions form a validation set. The loss function is evaluated on the validation set after each training epoch to monitor overtraining.

The results in Section 6.4 are produced by evaluating each NN on independent test MC datasets, which are not seen during training. The $t\bar{t}$ events used for training and testing originate from the same MC dataset but are separated to ensure an independent test of each NN’s performance.

6.2.4. Event selections

The simulated samples used to train the NNs do not have any specific event selections applied so as not to over-tune to specific event topologies. However, when testing the performance of each NN on different processes, event selections are applied to pick out a more refined final state that a typical analysis might use. These event selections are summarised in Table 6.4.

Table 6.4.: Summary of event selections applied when testing METNet performance.

Test data event selections	
Topology	Selections
$t\bar{t}$	$N_{\text{jet}} \geq 2$, $N_{b\text{-jet}} = 2$, $N_{\text{lepton}} \geq 1$
$WW \rightarrow \ell\nu\ell\nu$	2 OS signal/baseline leptons, $p_T^{\ell_1} > 25$ GeV, $p_T^{\ell_2} > 25$ GeV
$ZZ \rightarrow \nu\nu\ell\ell$	2 SFOS signal/baseline leptons, $p_T^{\ell_1} > 30$ GeV, $p_T^{\ell_2} > 25$ GeV, $ M_{\ell\ell} - M_Z < 15$ GeV
$Z \rightarrow \mu\mu$	2 OS signal/baseline muons, $p_T^{\mu_1} > 30$ GeV, $p_T^{\mu_2} > 25$ GeV, $ M_{\mu\mu} - M_Z < 15$ GeV
VBF $Z \rightarrow \mu\mu$	2 OS signal/baseline muons, $p_T^{\mu_1} > 30$ GeV, $p_T^{\mu_2} > 25$ GeV, $ M_{\mu\mu} - M_Z < 15$ GeV, $N_{\text{jet}} \geq 2$
VBF $Z \rightarrow ee$	2 OS signal/baseline electrons, $p_T^{e_1} > 30$ GeV, $p_T^{e_2} > 25$ GeV, $ M_{ee} - M_Z < 15$ GeV, $N_{\text{jet}} \geq 2$
VBF $W \rightarrow e\nu$	1 signal/baseline electron, $p_T^e > 25$ GeV, $N_{\text{jet}} \geq 2$
pMSSM model #1	None
2ℓ SUSY topology	2 OS signal/baseline muons, $p_T^{\ell_1} > 27$ GeV, $p_T^{\ell_2} > 9$ GeV, $ M_{\mu\mu} - M_Z > 15$ GeV, $N_{\text{jet}} \leq 1$, $N_{b\text{-jet}} = 0$,

When testing the NNs on $t\bar{t}$ events, leptonic decays of the top are preferred to obtain a final state with sources of real p_T^{miss} . Thus at least one signal lepton is required, along with at least two signal jets, exactly two of which are b -tagged. When testing on $WW \rightarrow \ell\nu\ell\nu$ events, two opposite sign (OS) signal leptons with $p_T > 25$ GeV are required, with a veto on the presence of any additional baseline leptons. The $W \rightarrow e\nu$ topology selection is similar, but requires a single electron with $p_T > 25$ GeV. When testing on $ZZ \rightarrow \nu\nu\ell\ell$, $Z \rightarrow \mu\mu$ and VBF $Z \rightarrow \ell\ell$ events, exactly two signal leptons are required, along with a veto on the presence of any additional baseline leptons. The two signal leptons must be the same flavour and oppositely signed (SFOS) and have an invariant mass within ± 15 GeV of the Z -boson mass. The highest- p_T lepton is set to be over the single-lepton-trigger efficiency plateau, by requiring $p_T > 30$ GeV. For the VBF topologies, an additional requirement of at least two signal jets is made.

The NNs are also tested on three pMSSM models taken from Chapter 4. The first, labelled ‘pMSSM #1’, has a large mass splitting between the $\tilde{\chi}_1^0$ and $\tilde{\chi}_2^0/\tilde{\chi}_1^\pm$, as shown in the mass spectrum in Figure 6.2. Therefore this model is used as an example bSM topology that produces large real p_T^{miss} . Since the pMSSM #1 event sample contains a




Figure 6.2.: Mass spectrum of pMSSM #1 from the EWKino scan in Chapter 4. The masses of the electroweakinos are shown and the dashed lines indicate the various decay modes. Produced using PySLHA [167].

variety of SUSY production mechanisms and decay modes, no specific event selections are required beyond the generator-level filtering applied in the production of the sample as outlined in Section 4.2.

Finally, to study the ability of the NNs to discriminate between topologies containing real and fake p_T^{miss} , a selection is applied to pick out a two-lepton final state SUSY topology. In this case, two pMSSM scenarios that are not excluded by the Run 2 searches in Chapter 4 are compared to a $Z \rightarrow \mu\mu$ and $WW \rightarrow \ell\nu\ell\nu$ background using selections inspired by the ATLAS two-lepton plus zero jet search in Reference [60]. This search was discussed previously in both Chapters 4 and 5. The pMSSM models are labelled ‘pMSSM #2’ and ‘pMSSM #3’. Exactly two signal muons are required as only the $Z \rightarrow \mu\mu$ background is considered (not $Z \rightarrow ee$), along with a veto on the presence of additional baseline leptons. The highest p_T muon is set to be over the single-muon-trigger efficiency plateau, with a small loss of signal acceptance, by requiring $p_T > 27$ GeV. There must be no b -tagged jets and a maximum of one light-flavour signal jet. The two muons must be oppositely charged and have an invariant mass at least ± 15 GeV away from the Z -boson mass.

6.3. Neural network architecture

The architectures used for each of the neural networks are presented in this section. First, the input features used are outlined before the architectural features common to both METNet_R and METNet_C are introduced. Then the particulars of the METNet_R and METNet_C set-ups will be presented.

6.3.1. Input features

The input features available to each NN are listed in Table 6.5 and include:

1. p_T^{miss} predictions and unique jet- and soft- terms for each working point. These provide each NN with information about each working point.
2. Lepton p_T^{miss} terms (note these are identical for every working point). For METNet_R this helps the NN build its own p_T^{miss} sum.
3. Additional variables which characterise the pile-up conditions and topology of each event. The NN can use this information to decide which working point is ‘best’ for a given event.

METNet_R uses all of these input features, while METNet_C uses a subset of 46 of these features as follows:

- Lepton p_T^{miss} term input features are dropped as they are not useful for identifying the best working point (as they are identical for each working point).
- The $p_{x,y}^{\text{miss}}$ and $p_{x,y}^{\text{miss, jet/soft}}$ features for each working point are dropped as the relevant information is already contained in the magnitudes p_T^{miss} and $p_T^{\text{miss, jet/soft}}$.

6.3.2. Common architecture

Both iterations of METNet use a multi-layer perceptron network consisting of an input and output layer, plus two fully connected hidden layers, each with 500 nodes. Each

Table 6.5.: Input features available to METNet. METNet_R uses all these features while METNet_C does not use the x or y components of the $p_{x,y}^{\text{miss}}$ input features or lepton terms.

Input features used for each working point			
p_T^{miss}	p_x^{miss}	p_y^{miss}	$\sum p_T$
$p_T^{\text{miss, jet}}$	$p_x^{\text{miss, jet}}$	$p_y^{\text{miss, jet}}$	$\sum p_T^{\text{jet}}$
$p_T^{\text{miss, soft}}$	$p_x^{\text{miss, soft}}$	$p_y^{\text{miss, soft}}$	$\sum p_T^{\text{soft}}$
Lepton p_T^{miss} term input features			
$p_T^{\text{miss, } e}$	$p_x^{\text{miss, } e}$	$p_y^{\text{miss, } e}$	$\sum p_T^e$
$p_T^{\text{miss, } \mu}$	$p_x^{\text{miss, } \mu}$	$p_y^{\text{miss, } \mu}$	$\sum p_T^\mu$
Additional input features			
$\langle \mu \rangle$	Mean number of interactions per bunch crossing		
μ	Actual number of interactions per bunch crossing		
N_{PV}	Number of primary vertices		
N_{PV^4}	Number of primary vertices with at least 4 associated tracks		
N_{trk}	Number of ID tracks associated with the primary vertex		
N_e^b	Number of baseline electrons		
N_μ^b	Number of baseline muons		
N_J^{Tight}	Number of jets satisfying the <i>Tight</i> working point selection		
N_J^{fwd}	Number of jets satisfying the <i>Loose</i> working point selection with $ \eta \geq 2.5$		
$\sum p_T^{\text{Fwd jets}}$	Sum of the p_T of jets satisfying the <i>Loose</i> working point selection with $ \eta \geq 2.5$		

hidden layer uses the Scaled Exponential Linear Unit (SELU) activation function [210],

$$\text{selu}(x) = \begin{cases} sx & , \quad x \geq 0 \\ s\alpha(e^x - 1) & , \quad x < 0 \end{cases} \quad (6.1)$$

where α and s are left as their default values $s = 1.05070098, \alpha = 1.67326324$. A layer normalisation layer (LayerNorm) [211] is also used after each hidden layer. This normalises the outputs of the previous layer for each training event and has been found to enable faster training.

Additionally, a normalisation pre-processing layer is used to standardise each input feature by subtracting the mean and dividing by the standard deviation. This is standard practice for deep learning problems as it ensures the features have a common scale. The mean and standard deviation are calculated using a 100,000 event sub-sample of the training data.

The hyperparameters used for both networks are summarised in Table 6.6. These hyperparameters have been selected through a combination of hyperparameter scans and known best-practice for regression problems, though there may be scope for further hyperparameter optimisation. Training is performed with Keras [212] using the TensorFlow backend [213]. Training proceeds for 25 epochs or until the loss of the validation set does not improve for 5 consecutive rounds (this is called *early stopping*).

Table 6.6.: Neural network hyperparameters.

Hyperparameter	Value
Optimiser	Adam [126]
Weight initialization	Kaiming He [214]
Learning rate	0.001
Batch size	256
Num. epochs	25
Early stopping	5
Number of hidden layers	2
Nodes per hidden layer	500

6.3.3. METNet_R architecture

METNet_R uses all 66 of the features in Table 6.5 as inputs. In addition to a pre-processing step to normalise the input features described previously, two additional pre-processing steps are applied for METNet_R:

- (1) Rotate each event such that $p_T^{\text{miss}, \text{Tight}}$ points along the x -axis by construction. This removes ϕ invariance from the inputs, increasing the statistical power of the training data.
- (2) Normalise each target variable by subtracting the mean and dividing by the standard deviation. The mean and standard deviation are calculated using a 100,000 event sub-sample of the training data.

The NN's outputs are transformed by the inverse of these steps to produce the final METNet_R prediction.




Figure 6.3.: Neural network architecture for METNet_R. Each hidden layer contains 500 nodes. Both input and target features are normalised.

The target variables used to train the network are $(p_x^{\text{miss, True}}, p_y^{\text{miss, True}})$. Alternative target variables have also been tested, including $\left(\frac{p_T^{\text{miss, Tight}}}{p_T^{\text{miss, True}}}, p_\phi^{\text{miss, True}}\right)$ and $(p_T^{\text{miss, True}}, p_\phi^{\text{miss, True}})$. Targeting $(p_x^{\text{miss}}, p_y^{\text{miss}})$ results in the best-performing network overall. In particular, targeting $p_\phi^{\text{miss, True}}$ causes issues due to the discontinuity at $\phi = \pm \pi$. The NN is designed to predict both the x and y components rather than just the magnitude p_T^{miss} , such that it could be used to calculate variables such as the transverse mass m_T and stransverse mass m_{T2} [95, 96] which are often used in searches for supersymmetry.

The METNet_R architecture consists of an input layer with 94 nodes (one for each input feature) followed by pre-processing layers to perform the event rotation and normalisation. This is followed by two hidden layers, each using the SELU activation and containing 500 nodes. A LayerNorm layer follows each hidden layer. The output layer uses a linear activation function, and the rotation and normalisation are reversed to produce the final outputs. This architecture is summarised in Figure 6.3.

The network's trainable parameters are determined during training by minimising the Huber loss function [127], $\mathcal{L}_{\text{Huber}}$, with the δ hyperparameter set to the Keras default value $\delta = 1.5$. A separate loss function is evaluated for each of the x - and y -components of $\mathbf{p}_T^{\text{miss}}$ which are then averaged to obtain an overall loss.

To reduce a bias in the NN's predictions towards the bulk of the $p_T^{\text{miss, True}}$ distribution of the training data, sample weights are used when calculating $\mathcal{L}_{\text{Huber}}$. The following procedure is used to calculate the weight for each training event. First, the $p_T^{\text{miss, True}}$ histogram is calculated between 0 and 300 GeV in 10 GeV width bins for a 100,000 event

sub-set of the training data. Then, for each event, the event weight is calculated as

$$w_e = \begin{cases} \frac{h_{\text{max.}}}{h_{50}} & , \quad p_T^{\text{miss, True}} < 50 \text{ GeV} \\ \frac{h_{\text{max.}}}{h_e} & , \quad 50 \leq p_T^{\text{miss, True}} \leq 300 \text{ GeV} \\ \frac{h_{\text{max.}}}{h_{300}} & , \quad p_T^{\text{miss, True}} > 300 \text{ GeV} \end{cases} \quad (6.2)$$

where $h_{\text{max.}}$ is the height of the most populated bin, h_e the height of the bin to which the event belongs, h_{50} is the height of the (40 GeV, 50 GeV) bin, and h_{300} is the height of the (300 GeV, 310 GeV) bin. This procedure effectively flattens the $p_T^{\text{miss, True}}$ histogram between 50 GeV and 300 GeV. An upper cut-off of 300 GeV is used to avoid huge event weights for lowly populated bins. A lower cut-off of 50 GeV reduced an observed bias towards zero from highly weighted training events with $p_T^{\text{miss, True}} = 0$. The calculated weights are then scaled such that the mean weight for all samples is equal to 1.

6.3.4. METNet_C architecture

The goal of METNet_C is to identify the optimal working point for each event. This requires a different architecture to METNet_R which performed a regression of $p_{x,y}^{\text{miss, True}}$.

First, each event must be assigned a class. This is taken to be the working point with the smallest squared error:

$$\text{Class} = \arg \min_{\text{WP}} \Delta_{\text{WP}} = \arg \min_{\text{WP}} \left[(p_x^{\text{miss, WP}} - p_x^{\text{miss, True}})^2 + (p_y^{\text{miss, WP}} - p_y^{\text{miss, True}})^2 \right]. \quad (6.3)$$

Other distance metrics could be used, but a Euclidean distance is picked for consistency with the root-mean-squared error (RMSE) calculation that is usually used to evaluate p_T^{miss} resolution in ATLAS. Note that in cases of a tie between multiple working points, a single working point is picked using the following order of priority: *Tight*, *Loose*, *Tenacious*, *TightClus*.

As is standard for multi-class classification problems, the categorical cross-entropy loss function is used, and the softmax activation function is used in the output layer. The output classification scores are used to form a weighted sum of the working points




Figure 6.4.: Neural network architecture for METNet_C. Each hidden layer contains 500 nodes.

to produce a final $p_{x,y}^{\text{miss}}$ prediction

$$p_{x,y}^{\text{miss, METNet}_C} = \sum_{\text{WP}} P_{\text{WP}} \cdot p_{x,y}^{\text{miss, WP}} \quad (6.4)$$

where P_{WP} is the classification score output by the network for working point WP. One could also simply use the working point with the highest output score as the final $p_{x,y}^{\text{miss}}$. However, since the network is not a perfect classifier, forming a weighted sum is found to produce the best performance.

Only the *Loose*, *Tight*, *TightClus* and *Tenacious* working points are used as target classes. The *Tighter* and *TightfJVT* working points are found to be too similar to *Tight* for the classifier to distinguish them. Additionally, the *Track* working point in general has a significant negative bias as it ignores neutral objects entirely. Therefore including it in the sum in Equation 6.4 causes a significant negative bias in the resulting p_T^{miss} variable, so the *Track* working point is excluded as a target for METNet_C.

Unlike METNet_R event weights are not used for METNet_C as $p_T^{\text{miss, True}}$ is not used as a training target so the network performance does not suffer from the same bias towards the p_T^{miss} bulk of the training set. Additionally, the x and y components of the input features are not used (as outlined in Section 6.3.1) so the event rotation pre-processing step is redundant and therefore switched off. The target scaling pre-processing step is also not required for METNet_C.

The METNet_C architecture consists of an input layer with 46 nodes (one for each input feature) followed by a pre-processing layer to normalise the input features, and a

block of two hidden layers with 500 nodes and two LayerNorm layers. This architecture is summarised in Figure 6.4.

6.3.5. Summary

A summary of the differences between the METNet_R and METNet_C architectures is given in Table 6.7.

Table 6.7.: Summary of the different architectures used for METNet_R and METNet_C.

	METNet _R	METNet _C
Target	$p_{x,y}^{\text{miss}}$	$\arg \min_{\text{WP}} \Delta_{\text{WP}}$
Loss function	$\mathcal{L}_{\text{Huber}}$	$\mathcal{L}_{\text{CCE}}$
Inputs	94 (Table 6.5)	46 (No lepton terms or x/y components)
Output layer activation function	Linear	Softmax
Pre-processing	Inputs normalised Targets normalised Rotate events	Inputs normalised
Sample weights	Flatten p_T^{miss} histogram	None
Trainable parameters	301,002	278,004

6.4. Results: Comparison with current p_T^{miss} working points

6.4.1. Resolution

Figure 6.5 shows the root-mean-square error (RMSE) for each working point and NN, where the ‘error’ is defined as the difference between the predicted value of $p_{x,y}^{\text{miss}}$ and $p_{x,y}^{\text{miss, True}}$. Both the x and y components enter this calculation, and their contributions are similar due to the rotational symmetry of the detector. The RMSE quantifies the




Figure 6.5.: Root-mean-square of the difference between the predicted value of $p_{x,y}^{\text{miss}}$ and $p_{x,y}^{\text{miss, True}}$ in bins of $p_T^{\text{miss, True}}$. METNet_R and METNet_C are shown along with current p_T^{miss} working points for (a) $t\bar{t}$ events, (b) VBF $W \rightarrow e\nu$, events. The lower panel shows the ratio with respect to the *Tight* working point, and the hatched band indicates the statistical uncertainty for the *Tight* working point.

resolution of the p_T^{miss} prediction – a smaller value indicates a more accurate prediction. The RMSE is shown binned in $p_T^{\text{miss, True}}$ for (a) $t\bar{t}$ and (b) VBF $W \rightarrow e\nu$ test MC events.

For $t\bar{t}$ events, both METNet_R and METNet_C have superior resolution to the p_T^{miss} working points up to $p_T^{\text{miss, True}} \approx 350$ GeV. Beyond this point, the performance of METNet_R deteriorates as the training statistics at higher $p_T^{\text{miss, True}}$ are diminished, and the event weight used saturates at $p_T^{\text{miss, True}} = 300$ GeV. This is an unavoidable issue with the METNet_R architecture – the $p_T^{\text{miss, True}}$ range of the training data will always be limited. The METNet_C performance is not constrained in this way, however, and it shows a consistent improvement on the other working points across the entire range of $p_T^{\text{miss, True}}$.

The training dataset is dominated by $t\bar{t}$, events, so the good performance for test $t\bar{t}$ events in the $p_T^{\text{miss, True}}$ region used in training is unsurprising. However, the training dataset did not include the VBF $W \rightarrow e\nu$ event topology. Nonetheless, both networks show the ability to extrapolate and produce a similar performance as seen for $t\bar{t}$ events.

Figure 6.6 also shows the RMSE, but in bins of number of primary vertices, for (a) $t\bar{t}$ events and (b) VBF $Z \rightarrow \mu\mu$ events. For both topologies METNet_R and METNet_C are more stable against pile-up than the current working points. The VBF $Z \rightarrow \mu\mu$ topology




Figure 6.6.: Root-mean-square of the difference between the predicted value of $p_{x,y}^{\text{miss}}$ and $p_{x,y}^{\text{miss, True}}$ in bins of number of primary vertices. METNet_R and METNet_C are shown along with current p_T^{miss} working points for (a) $t\bar{t}$ events, (b) VBF $Z \rightarrow \mu\mu$ events. The lower panel shows the ratio with respect to the *Tight* working point, and the hatched band indicates the statistical uncertainty for the *Tight* working point.

is not shown to the NNs during training and also contains no real p_T^{miss} . METNet_R and METNet_C generalise well to this new topology, with improved RMS compared to current p_T^{miss} working points.

6.4.2. Bias and tails

Ideally, for a given value of $p_T^{\text{miss, True}}$, the reconstructed p_T^{miss} should not show a bias towards higher or lower values. Bias can be studied through plots of the predicted p_T^{miss} distribution (Figure 6.7) and the response $p_T^{\text{miss}}/p_T^{\text{miss, True}}$ (Figure 6.8).

The METNet_C p_T^{miss} distribution has significantly improved agreement with $p_T^{\text{miss, True}}$ compared to the other working points for $t\bar{t}$ events (Figure 6.7a). However, METNet_R has much less stable performance, showing better agreement than the current working points in some bins but worse in others. For $Z \rightarrow \mu\mu$, shown in Figure 6.7b, both networks show a slight reduction in the fake- p_T^{miss} tail compared to the other working points.

Figure 6.8 shows the response $(p_T^{\text{miss}}/p_T^{\text{miss, True}})$ as a function of $p_T^{\text{miss, True}}$ for (a) $t\bar{t}$ and (b) $WW \rightarrow \ell\nu\ell\nu$ events. For METNet_R, the response dips below one around




Figure 6.7.: p_T^{miss} distribution for METNet_R, METNet_C, $p_T^{\text{miss, True}}$ and four p_T^{miss} working points, for (a) $t\bar{t}$ and (b) $Z \rightarrow \mu\mu$ events. Events are unweighted.




Figure 6.8.: Response of the predicted p_T^{miss} for METNet_R, METNet_C, and four current p_T^{miss} working points (a) $t\bar{t}$, (b) $WW \rightarrow \ell\nu\ell\nu$ events.

$p_T^{\text{miss, True}} \approx 80$ GeV, which indicates a negative bias. As with other metrics shown, the response also deteriorates for METNet_R beyond around $p_T^{\text{miss, True}} \approx 300$ GeV. METNet_C also shows a small negative bias, most notably for $t\bar{t}$ events, but otherwise has good closure for the range of values shown. Note that in all cases the response rises as $p_T^{\text{miss, True}}$ approaches 0 by the way the variable is constructed.




Figure 6.9.: Two p_T^{miss} performance plots for several p_T^{miss} working points, for a pMSSM signal sample that produces large real p_T^{miss} : (a) Unweighted p_T^{miss} distribution, (b) The root-mean-square of the difference between the predicted value of $p_{x,y}^{\text{miss}}$ and $p_{x,y}^{\text{miss, True}}$ in bins of number of primary vertices.

It is noted that this inherent *negative* bias could benefit the performance for topologies containing no real p_T^{miss} , such as $Z \rightarrow \mu\mu$, for which the current working points have a *positive* bias. It is unclear how much of the improvement in the fake p_T^{miss} tails in Figure 6.7b is due to the negative bias vs. a genuine ability to extrapolate to this topology.

In Figure 6.9, the NNs are tested on pMSSM model #1 which produces large p_T^{miss} due to large mass-splittings between the LSP and other electroweakinos. METNet_R has very poor performance for this topology. As seen previously, it can't extrapolate to events with $p_T^{\text{miss, True}}$ beyond the training dataset, which can be seen in the premature drop-off in the p_T^{miss} distribution in Figure 6.9a. The RMSE in Figure 6.9b is also significantly worse for METNet_R than the other working points. METNet_C on the other hand does not suffer from these issues. It shows a small improvement in RMSE for $\langle \mu \rangle > 20$ and otherwise has comparable performance to the current working points.

6.4.3. Signal vs. background separation

Here the NN's are evaluated in the context of a toy search for supersymmetry in events with two leptons and zero jets, based on Reference [141]. Figure 6.10 compares the signal-to-background separation power of the *Loose* working point, the *Tenacious* working point, METNet_R, and METNet_C. The *Loose* and *Tenacious* working points are picked for




Figure 6.10.: Distributions of p_T^{miss} for two pMSSM signal points, and WW and $Z \rightarrow \mu\mu$ backgrounds, for (a) *Loose* p_T^{miss} , (b) *Tenacious* p_T^{miss} , (c) METNet_R and (d) METNet_C. Events are weighted to 139 fb^{-1} , corresponding to the 2015-2018 ATLAS proton-proton collision dataset luminosity. The lower plot shows signal significance calculated by applying a lower-bound cut at each x -axis bin value.

comparison as the worst and best performing working points for this event selection. The distribution of two pMSSM signal points (pMSSM #2 and pMSSM #3) are compared to two of the main Standard Model backgrounds: $WW \rightarrow \ell\nu\ell\nu$ and $Z \rightarrow \mu\mu$. The kinematic selection applied is the ‘2 ℓ SUSY topology’ from Table 6.4. The lower panel shows the signal significance, as defined in Reference [215], for each pMSSM point for a lower-bound selection at each value on the x -axis.

Table 6.8 shows the maximum signal significance achieved for each p_T^{miss} definition in Figure 6.10 plus the *Tight* and *TightClus* working points. Both METNet_R and METNet_C

Table 6.8.: Maximum signal significance for each working point in Figure 6.10 (plus *Tight* and *TightClus*).

p_T^{miss} definition	Max. signal significance	
	pMSSM #1	pMSSM #2
<i>Loose</i>	0.92	0.79
<i>Tight</i>	1.08	1.01
<i>TightClus</i>	1.23	1.09
<i>Tenacious</i>	1.34	1.06
METNet _R	1.47	1.18
METNet _C	1.63	1.26

Table 6.9.: Remaining events for different levels of $Z \rightarrow \mu\mu$ rejection using cuts on *Tight* p_T^{miss} (Tight), regressed METNet_R (Reg.) and the classifier METNet_C (Class.). Values are shown for three signals: $WW \rightarrow \ell\nu\ell\nu$, pMSSM #2 and pMSSM #3.

Process	Percentage of events after X% $Z \rightarrow \mu\mu$ rejection								
	X=90%			X=99%			X=99.9%		
	Tight	Reg.	Class.	Tight	Reg.	Class.	Tight	Reg.	Class.
$WW \rightarrow \ell\nu\ell\nu$	65%	63%	70%	24%	29%	34%	7%	10%	13%
pMSSM #1	86%	83%	88%	60%	65%	70%	34%	41%	49%
pMSSM #2	91%	92%	94%	75%	77%	80%	55%	60%	64%

produce a higher signal significance than the current working points, indicating their potential for use in searches for supersymmetry. METNet_C produces the best sensitivity overall.

As a quantitative comparison of the discrimination between real p_T^{miss} and fake p_T^{miss} topologies, Table 6.9 shows the percentage of $WW \rightarrow \ell\nu\ell\nu$, pMSSM #2 and pMSSM #3 events surviving p_T^{miss} selections which remove 90, 99 and 99.9% of $Z \rightarrow \mu\mu$ events, for the *Tight* working point, METNet_R, and METNet_C. METNet_C has the best separation power in all cases. METNet_R shows an improvement compared to *Tight* for 99% and 99.9% $Z \rightarrow \mu\mu$ rejection but does not separate real and fake p_T^{miss} as well as METNet_C.

6.4.4. Discussion: METNet_R vs. METNet_C

The performance of the predicted p_T^{miss} for two neural network architectures, METNet_R and METNet_C, has been presented compared to current working points. Overall, both METNet_R and METNet_C demonstrate some improvements in performance compared to current p_T^{miss} working points. However, the performance improvement for METNet_R is restricted to topologies and $p_T^{\text{miss, True}}$ ranges that are similar to the training dataset while METNet_C is far more robust.

The RMSE for METNet_R deteriorates severely for events with $p_T^{\text{miss, True}}$ larger than around 350 GeV. While this could be partially combatted by including slices of events with higher $p_T^{\text{miss, True}}$ in the training set, it is not possible to train the network on an infinite range of events. Therefore, the performance of the METNet_R architecture will always be constrained to a finite range in $p_T^{\text{miss, True}}$. METNet_C does not have this problem as, by construction, it can only perform as badly as the worst-performing working point.

Each NN is shown multiple event topologies during training. Consequently, they can generalise to a broad range of event topologies including those not seen in training. METNet_C shows a stronger ability to generalise, indicating that it has learned a topology-independent relationship between the optimal working point and the input features. METNet_R on the other hand can generalise to topologies with some similarity to those seen in training, such as $Z \rightarrow \mu\mu$ when $ZZ \rightarrow \nu\nu\ell\ell$ is seen during training. However completely novel topologies, like the pMSSM #1 signal point, result in significantly worse performance than the current working points.

It is concluded that a single METNet_C network could be used as a general tool for a wide variety of ATLAS searches and measurements, where it will produce either similar or improved performance to the existing working points. The usage of METNet_R on the other hand would need to be restricted to analyses targeting similar event topologies and $p_T^{\text{miss, True}}$ ranges to those used to train the NN, which is a significant disadvantage.

Event weights are used when training METNet_R to reduce an observed bias towards the bulk of the training data. Nonetheless, a negative bias remains for some topologies. METNet_R is optimised to minimise the mean-squared error (and mean-absolute error) through the Huber loss function. However, there is no requirement built into this loss function for an unbiased estimation. Similarly, the Huber loss function penalises the network for poor event-by-event error but is indifferent to the agreement of the predicted p_T^{miss} *distribution* with the true p_T^{miss} distribution. Therefore, the bias seen in

METNet_R's predictions is not surprising, as the NN is not punished during training for producing a biased prediction. During the development and optimisation of METNet_R, a distribution-matching contribution to the loss function was considered through the Sinkhorn function [216], which penalises the network for poor agreement in the predicted distribution for each batch. This somewhat improved the agreement of the predicted p_T^{miss} distribution with $p_T^{\text{miss, True}}$, but did not alleviate these issues and resulted in inferior resolution.

METNet_C shows a much clearer and more robust improvement in the agreement of the p_T^{miss} distribution with $p_T^{\text{miss, True}}$. Some response plots do show a small negative bias for some topologies, but this is less severe than for METNet_R.

The utility of METNet_R and METNet_C in searches for supersymmetry is also considered, and is particularly relevant to this thesis. In the toy analysis of two pMSSM signal samples compared to SM backgrounds in Section 6.4.3, both METNet_R and METNet_C showed improved signal vs. background and real p_T^{miss} vs. fake p_T^{miss} discrimination power compared to current working points. However, the performance of METNet_R on a pMSSM signal with large mass splittings and therefore large $p_T^{\text{miss, True}}$ is very poor due to the aforementioned issues with extrapolation beyond the training set. Therefore, METNet_R would be useless in any SUSY search targeting high p_T^{miss} bins – this would rule out many SUSY searches. METNet_C does not suffer from this problem and showed the best overall ability to discriminate signal and background topologies so would be a far more robust choice for a search.

Overall, by almost all metrics considered METNet_C is the superior architecture for a general p_T^{miss} definition that can offer improvements to both p_T^{miss} reconstruction performance and the discrimination power in searches for new physics. Historically, the development of METNet began with METNet_R and included almost two years of optimisation and attempts to address the issues with extrapolation highlighted in this section. Ultimately, my idea to switch the learning task to classification rather than regression solved these problems. METNet_R in principle had greater potential, as a neural network that could learn a general topology-independent map from reconstructed p_T^{miss} to $p_T^{\text{miss, True}}$ could be transformative to the performance of p_T^{miss} in ATLAS. It seems that, at least with the input features provided, the NN could not learn such a general map. METNet_C on the other hand has a more moderate offering. It cannot perform worse than the worst existing working point, but also cannot improve performance too far from the best working point. These constraints produce a far more robust p_T^{miss} prediction than METNet_R that still shows improvements in both p_T^{miss} resolution and sensitivity in

searches for new physics. Therefore I encourage you, the reader, to think carefully when developing a machine learning model about whether you have picked the best learning task for the problem at hand.

6.5. Classification scores

Given the strong performance of METNet_C, it is worth taking a closer look at the classification scores from which the final p_T^{miss} prediction is built. How well is the classifier distinguishing each working point for different topologies? The plots displayed in Figure 6.11 help to answer this question. The y -axis in each plot shows two metrics for each bin and working point: (i) the average output score for each working point, indicated by crosses, and (ii) the fraction of events for which each working point is the closest to truth, indicated by dots and referred to as the ‘true fraction’. For a classifier behaving as desired, the dots and crosses of each colour should align very closely. The plots are shown for $t\bar{t}$ and $Z \rightarrow \mu\mu$ events in bins of $\langle\mu\rangle$, N_{PV} and $\sum p_T$ where the sum is over all objects entering the *Tight* working point p_T^{miss} calculation.

For $t\bar{t}$ (the first column), the average classifier score agrees very closely with the true fraction for each working point and bin. The classifier has learned, for example, to output a high score for *TightClus* at very low pile-up (ie. low $\langle\mu\rangle$ and N_{PV}) but not at high pile-up.

For $Z \rightarrow \mu\mu$ (the second column) – a topology not seen by the classifier in training – the agreement is not as close, but nonetheless the general shape of the classifier predictions and true fraction agree. The classifier struggles in particular between *Tight* and *TightClus*, especially for lower pile-up. This is likely because at low pile-up the *Tight* and *TightClus* working points will produce very similar inputs to the network as they differ only in the soft term. Overall, while the classifier is somewhat over-tuned to $t\bar{t}$, the topology that dominated its training set, the map it has learned from the inputs to the classifier scores can generalise reasonably well to other topologies.

A similar pattern can be seen in the confusion matrices shown in Figure 6.12. Here each box is coloured and labelled by the % of events with true class given by the y -axis for which the highest scoring output class is that on the x -axis. A perfect classifier would have 100% and 0% in every diagonal and off-diagonal box, respectively. For $t\bar{t}$, there is some confusion between *Tight* and *TightClus*. However, the classifier very




Figure 6.11.: The average classifier scores and the fraction of events belonging to each class for $t\bar{t}$ and $Z \rightarrow \mu\mu$ simulated test events with 2018 conditions, in bins of (a) average interactions per bunch crossing, (b) number of primary vertices, and (c) sum of the p_T of all objects contributing to the $Tight\ p_T^{\text{miss}}$ working point.





Figure 6.12.: Confusion matrix for the classifier's prediction of the 'best' working point, for (a) $t\bar{t}$ and (b) $Z \rightarrow \mu\mu$ simulated test events with 2018 conditions.

rarely misclassifies a *Tight* or *TightClus* event as *Loose* or *Tenacious*. For $Z \rightarrow \mu\mu$, the confusion between *Tight* and *TightClus* is greater, with 77% of true *TightClus* events classified as *Tight*. In general, there is a bias towards the *Tight* working point for both $t\bar{t}$ and $Z \rightarrow \mu\mu$. This may be related to the procedure for assigning a ‘true’ class in the case of ties between multiple working points, which gives first priority to *Tight*.

6.6. Outlook

This chapter presents the performance of two machine-learning-based p_T^{miss} working points: METNet_R and METNet_C. Both provide improvements compared to current p_T^{miss} working points, including improved resolution and pile-up resilience for a range of topologies. METNet_C in particular shows potential for use in general physics analyses. However, there is room for further improvements, and some avenues for further study are outlined here.

There is scope to further optimise the hyperparameters used by the NNs presented here. A limited grid scan was performed early in the development of METNet to produce an architecture that suited the problem. However, the training setup has since been modified to include more training events, training sample topologies and input features. Thus the current choice of hyperparameters is unlikely to be optimal. Nonetheless, the current hyperparameters produce a NN with strong performance, and METNet_C as defined here is very competitive compared to existing working points.

For METNet_C, the training data used is not balanced in the target classes. The *Tight* working point in particular is over-represented in the training data. It may be beneficial to show METNet_C a more balanced training set through either class sampling or sample weights, to enhance its predictions for less common classes. Additionally, a priority order is used when assigning classes to events in the case of a tie: *Tight*, *Loose*, *Tenacious*, *TightClus*. This produces a bias towards the *Tight* working point. There may be a better way to handle the situation of a tie between multiple working points. For example, in the case of a tie the target class could be selected randomly to avoid this bias. Alternatively, every possible combination of working points that tie could be one-hot encoded to produce an output score for each.

Additionally, the topologies used in training could be expanded. Fake p_T^{miss} topologies such as $Z \rightarrow \mu\mu$ were not used in training as they induced a severe bias towards zero

in the predictions of METNet_R. However, METNet_C would not suffer from such a bias, and the inclusion of $Z \rightarrow \mu\mu$ (and $Z \rightarrow ee$), which favours the *Tenacious* working point, may improve the generalisation power of the NN.

One could also study a NN architecture which directly regresses the linear coefficients in the METNet_C sum (Equation 6.4). The choice of a linear sum of the METNet_C classifier outputs may also be suboptimal – there are in principle infinite ways to combine the p_T^{miss} working points using the classifier scores.

To develop METNet_C into a fully-fledged recommendation for use in ATLAS analyses, two additional studies are required. First, the agreement between data and MC should be evaluated. Given that METNet_C is simply a linear combination of working points which have good closure themselves (as shown in Reference [94] for example), it is expected that good agreement would be found. It remains to be seen whether METNet_C demonstrates improved data/MC agreement compared to current working points – this is possible if it is selecting the working point with the best agreement for each event, but it has not been directly optimised to do this. Secondly, systematic uncertainties must be derived for use in analysis. For METNet_C, the individual working points have their own associated systematic uncertainties which can be propagated in the final sum. However, it should be established whether an additional systematic uncertainty on the NN’s output coefficients is necessary and if so what method should be used.

The METNet_R architecture could be developed further by modelling $p_{x,y}^{\text{miss, True}}$ as a Gaussian distribution for each event with central value $p_{x,y}^{\text{miss, True}}$ and width $\sigma_{x,y}$. This is achieved through the Gaussian negative log-likelihood (GNLL) loss [217, 218] function:

$$\mathcal{L}_{\text{GNLL}} = \log \sigma + 0.5 \left(\frac{y - \hat{y}}{\sigma} \right)^2. \quad (6.5)$$

The resulting NN has 4 outputs: $(p_x^{\text{miss, NN}}, \sigma_x, p_y^{\text{miss, NN}}, \sigma_y)$. The predicted values of $\sigma_{x,y}$ can be treated as a ‘confidence’ in the NNs prediction and used to define a machine-learning-based p_T^{miss} significance variable

$$\text{ML } p_T^{\text{miss}} \text{ significance} = p_T^{\text{miss, NN}} / \sigma \quad (6.6)$$

where

$$p_T^{\text{miss, NN}} = \sqrt{(p_x^{\text{miss, NN}})^2 + (p_y^{\text{miss, NN}})^2} \quad (6.7)$$

and

$$\sigma = \frac{\sqrt{(p_x^{\text{miss, NN}} \sigma_x)^2 + (p_y^{\text{miss, NN}} \sigma_y)^2}}{p_T^{\text{miss, NN}}}. \quad (6.8)$$

This variable could discriminate between real and fake p_T^{miss} . Initial studies of such an architecture were performed (by the author of this thesis) in Reference [3] but are not included in this thesis. However this will inevitably be affected by the same issues as METNet_R, and be limited to similar topologies and $p_T^{\text{miss, True}}$ ranges as the training data. It could nonetheless be useful in an analysis struggling with backgrounds containing fake p_T^{miss} , if a targeted training was performed with the relevant production processes.

Finally, there exists in principle an infinite number of p_T^{miss} working points, each with different jet selection thresholds on variables such as jet p_T and JVT. The working points defined by ATLAS are a discrete set of benchmark points from the continuum of possible choices. The studies presented here show that utilising different jet selections for different topological and pile-up conditions is beneficial. This raises two questions. First, are there worthwhile gains to be made from defining additional working points to be used in METNet_C's combination? For example, there may be choices of jet selection and soft term that are only optimal for a peculiar fraction of ATLAS events and thus not useful for general analyses but could be deployed where appropriate by METNet_C. Second, could a machine learning architecture be designed to decide the jet selection itself, rather than picking from a limited pool of pre-defined working points? The relationship between the jet selection used and the final p_T^{miss} is complex, as it defines which low-level objects enter either jet reconstruction or the soft term. Such an architecture would most likely need access to these low-level inputs, and it is not clear whether such an approach is feasible.

To conclude, machine learning techniques can significantly improve the agreement between reconstructed and true p_T^{miss} . The results shown in this chapter strongly motivate the use of METNet_C as a recommended working point in ATLAS physics analyses. With further work to derive systematic uncertainties and study the impact on data/MC agreement, METNet_C could be a robust choice for p_T^{miss} in Run 3 analyses, including searches for supersymmetry. The studies presented here also suggest that the improvements from this method could enhance the sensitivity of searches for supersymmetry. While the gain from METNet_C alone may be incremental, the consolidation of similar performance enhancements across the suite of ATLAS analysis objects will provide important improvements to the sensitivity of ATLAS to new physics during Run 3 and beyond.

Chapter 7.

Summary and outlook

This thesis presents three approaches to enhancing the sensitivity of the ATLAS experiment to the electroweak production of supersymmetric particles.

In Chapter 4 a suite of Run 2 searches are interpreted in the context of the pMSSM. This captures a richer phenomenology than the original simplified model interpretations, producing a clearer picture of the resulting global constraints on supersymmetric scenarios. Importantly, it reveals many scenarios with mass-spectra that nominally lie within published exclusion contours that remain viable. This enhanced understanding of ATLAS' sensitivity to electroweak supersymmetry provides a legacy of the Run 2 results and will guide the direction of new searches performed with Run 3 ATLAS data.

Then, in Chapter 5, a statistical combination of three searches for the pair-production of charginos is performed. This results in a world-leading exclusion limit on chargino masses in simplified scenarios.

Finally, in Chapter 6, a machine-learning-based estimation of missing transverse momentum, named METNet, is presented. This offers significantly improved resolution compared to the missing transverse momentum estimates currently used by ATLAS analyses. A moderate enhancement in the separation between pMSSM signal and Standard Model background processes is seen when using METNet. This, and other improvements to ATLAS reconstruction performance, are complementary to the sensitivity enhancements that can be gained from analysis selection optimisation and statistical combinations.

Where does this leave the hunt for electroweak supersymmetry? The pMSSM analysis in Chapter 4 produced several benchmark scenarios with non-simplified spectra which are not excluded by previous ATLAS searches or other constraints. It is vital to target these

with analyses of Run 3 data. Hopefully, these analyses will make use of METNet to access as accurate a p_T^{miss} estimate as possible. As demonstrated in Chapter 5, a statistical combination of multiple orthogonal analyses is a powerful method of maximising the sensitivity of searches. Future searches should take heed of this and, as was done in Run 2, build orthogonality between analysis selections by construction with a view to future combinations.

The pMSSM model sets from Chapter 5 can be used for the interpretation of future searches, such that the unique sensitivity offered by each new analysis can be evaluated. In this context, there may be scope to develop the pMSSM workflow outlined in Section 4.2 to allow for the continuous integration of pMSSM interpretations with developing SUSY analyses in ATLAS. This could provide valuable final results demonstrating the sensitivity of a search to realistic pMSSM scenarios, and also ensure a developing analysis does not over-tune to simplified models.

To conclude, despite the optimistic (or pessimistic, depending on your point of view) exclusion limits reported by Run 2 searches for simplified models of SUSY, vast swathes of parameter space remain uncovered. Run 3 of the LHC is now underway, and the size of the ATLAS dataset will double in size over the next few years. Experimentalists must use every tool in their arsenal to target viable SUSY scenarios in pursuit of a new discovery. By performing targeted searches optimised to viable non-simplified scenarios, making use of improvements to reconstruction performance such as METNet, and performing statistical combinations of multiple searches, we can push ATLAS to the limit of what it can reveal about the fundamental building blocks of our universe.

Appendix A.

Extra combination plots

Figures A.1, A.2 and A.3 show pull plots for the observed combined exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (300, 100)$, $(600, 200)$ and $(900, 300)$ points respectively. A different set of analyses are combined for each of these points, as indicated in the captions. Only systematic uncertainties for which the pull for the combined fit, rounded to two decimal places, is different compared to one of the original analyses are displayed. The original analysis pull is displayed in black and the combination pull in red.

Several pulls can be seen for the combination and individual analyses, however these mostly have a pull of similar magnitude to the individual analyses. The most significant differences are for muon-related systematic uncertainties for 1L and 2L0J. However, all pulls remain within 1σ so are not considered problematic.

Figure A.4 shows the correlations between systematic uncertainties for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (600, 200)$ point.




Figure A.1.: Pull plot from an observed exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (300, 100)$ mass point, which combines the 2L0J and 1L analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pull is shown in red, the original analysis pull in black.




Figure A.2.: Pull plot from an observed exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (600, 200)$ mass point, which combines the 2L0J, 1L and FullHad analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pull is shown in red, the original analysis pull in black.




Figure A.3.: Pull plot from an observed exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (900, 300)$ mass point, which combines the 1L and FullHad analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pull is shown in red, the original analysis pull in black.




Figure A.4.: Correlations between systematic uncertainties for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (600, 200)$ mass point, for which the 2L0J wave 1, 1L and FullHad analyses are combined.

Colophon

This thesis was made in $\text{\LaTeX} 2_{\varepsilon}$ using the “hepthesis” class [\[219\]](#).

Bibliography

- [1] ATLAS Collaboration, *ATLAS Run 2 searches for electroweak production of supersymmetric particles interpreted within the p MSSM*, Tech. Rep. ATLAS-CONF-2023-055, CERN, Geneva, 2023.
<https://cds.cern.ch/record/2870222>.
- [2] ATLAS Collaboration, *Run 2 results of searches for charginos and neutralinos at the ATLAS experiment using a statistical combination*, Tech. Rep. ATLAS-CONF-2023-046, CERN, Geneva, 2023.
<https://cds.cern.ch/record/2870114>.
- [3] ATLAS Collaboration, *METNet: A combined pT_{miss} working point using a neural network with the ATLAS detector*, Tech. Rep. ATL-PHYS-PUB-2021-025, CERN, Geneva, 2021. <https://cds.cern.ch/record/2776653>.
- [4] B. Hodkinson, *METNet: A combined missing transverse momentum working point using a neural network with the ATLAS detector*, [PoS EPS-HEP2021 \(2022\) 625](#).
- [5] B. Hodkinson, *Improving ATLAS Hadronic Object Performance with ML/AI Algorithms*, [arXiv:2307.01388 \[hep-ex\]](#).
- [6] M. Leigh, *Analysis of a deep neural network for missing transverse momentum reconstruction in ATLAS*, Master's thesis.
<http://hdl.handle.net/11427/32401>.
- [7] E. Noether, *Invariant variation problems*, [Transport Theory and Statistical Physics](#) **1** no. 3, (1971) 186–207.
- [8] S. L. Glashow, *Partial-symmetries of weak interactions*, *Nuclear Physics* **22** no. 4, (1961) 579.
- [9] S. Weinberg, *A model of leptons*, *Phys. Rev. Lett.* **19** no. 21, (1967) 1264.
- [10] A. Salam, *Weak and electromagnetic interactions*, *Conf. Proc.* **C680519** (1968)

- 367.
- [11] J. Liu, *Discovery strategies for dark matter and Higgsinos at the LHC*, 2019.
<https://cds.cern.ch/record/2683899>. Presented 29 May 2019.
 - [12] A. Hook, *TASI Lectures on the Strong CP Problem and Axions*, 2021.
 - [13] Particle Data Group Collaboration, R. L. Workman et al., *Review of Particle Physics*, **PTEP** **2022** (2022) 083C01.
 - [14] Super-Kamiokande Collaboration, *Evidence for Oscillation of Atmospheric Neutrinos*, **Phys. Rev. Lett.** **81** (1998) 1562–1567.
 - [15] SNO, *Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory*, **Physical Review Letters** **89** no. 1, (2002).
 - [16] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, **Phys. Rev. Lett.** **10** (1963) 531–533.
 - [17] M. Kobayashi and T. Maskawa, *CP-Violation in the Renormalizable Theory of Weak Interaction*, **Progress of Theoretical Physics** **49** no. 2, (1973) 652–657.
 - [18] Z. Maki, M. Nakagawa, and S. Sakata, *Remarks on the Unified Model of Elementary Particles*, **Progress of Theoretical Physics** **28** no. 5, (1962) 870–880.
 - [19] B. Pontecorvo, *Inverse beta processes and nonconservation of lepton charge*, *Zh. Eksp. Teor. Fiz.* **34** (1957) 247.
 - [20] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, **Phys. Lett. B** **716** (2012) 1, [arXiv:1207.7214 \[hep-ex\]](#).
 - [21] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, **Phys. Lett. B** **716** (2012) 30, [arXiv:1207.7235 \[hep-ex\]](#).
 - [22] Muon g-2 Collaboration Collaboration, Muon g-2 Collaboration, *Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm*, **Phys. Rev. Lett.** **126** (2021) 141801.
 - [23] Muon g-2 Collaboration Collaboration, Muon g-2 Collaboration, *Final report of the E821 muon anomalous magnetic moment measurement at BNL*, **Phys. Rev. D**

- 73** (2006) 072003.
- [24] E. Graverini, *Flavour anomalies: a review*, **Journal of Physics: Conference Series** **1137** (2019) 012025.
- [25] J. C. Kapteyn, *First Attempt at a Theory of the Arrangement and Motion of the Sidereal System*, **Astrophys. J.** **55** (1922) 302–328.
- [26] V. C. Rubin and W. K. Ford, Jr, *Rotation of the andromeda nebula from a spectroscopic survey of emission regions*, **Astrophys. J.** **159**: 379-403(Feb 1970).
- [27] E. Corbelli and P. Salucci, *The extended rotation curve and the dark matter halo of M33*, **Monthly Notices of the Royal Astronomical Society** **311** no. 2, (2000) 441–447.
- [28] J. A. Tyson, G. P. Kochanski, and I. P. D. Antonio, *Detailed Mass Map of CL 0024+1654 from Strong Lensing*, **The Astrophysical Journal** **498** no. 2, (1998) L107–L110.
- [29] A. Refregier, *Weak Gravitational Lensing by Large-Scale Structure*, **Annual Review of Astronomy and Astrophysics** **41** no. 1, (2003) 645–668.
- [30] A. D. Lewis, D. A. Buote, and J. T. Stocke, *Chandra Observations of A2029: The Dark Matter Profile Down to below $0.01r_{vir}$ in an Unusually Relaxed Cluster*, **The Astrophysical Journal** **586** no. 1, (2003) 135–142.
- [31] S. W. Allen, A. E. Evrard, and A. B. Mantz, *Cosmological Parameters from Observations of Galaxy Clusters*, **Annual Review of Astronomy and Astrophysics** **49** no. 1, (2011) 409–470.
- [32] M. Tegmark, M. A. Strauss, M. R. Blanton, et al., *Cosmological parameters from SDSS and WMAP*, **Phys. Rev. D** **69** (2004) 103501.
- [33] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, et al., *Planck 2018 results. VI. Cosmological parameters.*, **Astronomy & Astrophysics** **641** (2020) A6.
- [34] M. Lisanti, *Lectures on Dark Matter Physics*, in *New Frontiers in Fields and Strings*. WORLD SCIENTIFIC, Nov, 2016.
- [35] S. P. Martin, *A Supersymmetry Primer*, **Adv. Ser. Direct. High Energy Phys.** **18** (1998) 1, [arXiv:hep-ph/9709356](#).
- [36] I. J. R. Aitchison, *Supersymmetry and the MSSM: An Elementary Introduction*,

- 2005.
- [37] Y. A. Golfand and E. P. Likhtman, *Extension of the Algebra of Poincare Group Generators and Violation of p Invariance*, JETP Lett. **13** (1971) 323.
 - [38] D. V. Volkov and V. P. Akulov, *Is the Neutrino a Goldstone Particle?*, Phys. Lett. B **46** (1973) 109.
 - [39] J. Wess and B. Zumino, *Supergauge Transformations in Four-Dimensions*, Nucl. Phys. B **70** (1974) 39.
 - [40] J. Wess and B. Zumino, *Supergauge Invariant Extension of Quantum Electrodynamics*, Nucl. Phys. B **78** (1974) 1.
 - [41] S. Ferrara and B. Zumino, *Supergauge Invariant Yang-Mills Theories*, Nucl. Phys. B **79** (1974) 413.
 - [42] A. Salam and J. A. Strathdee, *Supersymmetry and Nonabelian Gauges*, Phys. Lett. B **51** (1974) 353.
 - [43] N. Sakai, *Naturalness in supersymmetric GUTS*, Z. Phys. C **11** (1981) 153.
 - [44] R. R. S. Dimopoulos and F. Wilczek, *Supersymmetry and the scale of unification*, Phys. Rev. D **24** (1981) 1681.
 - [45] L. E. Ibanez and G. G. Ross, *Low-Energy predictions in supersymmetric grand unified theories*, Phys. Lett. B **105** (1981) 439.
 - [46] S. Dimopoulos and H. Georgi, *Softly broken supersymmetry and $SU(5)$* , Nucl. Phys. B **193** (1981) 150.
 - [47] Super-Kamiokande Collaboration, *Search for Proton Decay via $p \rightarrow e^+ \pi^0$ and $p \rightarrow \mu^+ \pi^0$ in a Large Water Cherenkov Detector*, Phys. Rev. Lett. **102** (2009) 141801.
 - [48] Super-Kamiokande Collaboration, *Search for nucleon decay into charged antilepton plus meson in Super-Kamiokande I and II*, Physical Review D **85** no. 11, (2012).
 - [49] G. R. Farrar and P. Fayet, *Phenomenology of the Production, Decay, and Detection of New Hadronic States Associated with Supersymmetry*, Phys. Lett. B **76** (1978) 575.
 - [50] R. N. Mohapatra, *Supersymmetry and R -parity: an Overview*, 1503.

- [51] A. H. Chamseddine, R. L. Arnowitt, and P. Nath, *Locally Supersymmetric Grand Unification*, Phys. Rev. Lett. **49** (1982) 970.
- [52] R. Barbieri, S. Ferrara, and C. A. Savoy, *Gauge models with spontaneously broken local supersymmetry*, Phys. Lett. B **119** (1982) 343.
- [53] G. L. Kane, C. F. Kolda, L. Roszkowski, and J. D. Wells, *Study of constrained minimal supersymmetry*, Phys. Rev. D **49** (1994) 6173.
- [54] M. Dine and W. Fischler, *A phenomenological model of particle physics based on supersymmetry*, Phys. Lett. B **110** (1982) 227.
- [55] L. Alvarez-Gaume, M. Claudson, and M. B. Wise, *Low-energy supersymmetry*, Nucl. Phys. B **207** (1982) 96.
- [56] C. R. Nappi and B. A. Ovrut, *Supersymmetric extension of the $SU(3) \times SU(2) \times U(1)$ model*, Phys. Lett. B **113** (1982) 175.
- [57] ATLAS Collaboration, *Search for new phenomena in final states with large jet multiplicities and missing transverse momentum using $\sqrt{s} = 13$ TeV proton–proton collisions recorded by ATLAS in Run 2 of the LHC*, JHEP **10** (2020) 062, [arXiv:2008.06032 \[hep-ex\]](#).
- [58] ATLAS Collaboration, *Search for squarks and gluinos in final states with jets and missing transverse momentum using 139 fb^{-1} of $\sqrt{s} = 13$ TeV pp collision data with the ATLAS detector*, JHEP **02** (2021) 143, [arXiv:2010.14293 \[hep-ex\]](#).
- [59] ATLAS Collaboration, *Search for squarks and gluinos in final states with one isolated lepton, jets, and missing transverse momentum at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Eur. Phys. J. C **81** (2021) 600, [arXiv:2101.01629 \[hep-ex\]](#).
- [60] ATLAS Collaboration, *Search for electroweak production of charginos and sleptons decaying into final states with two leptons and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions using the ATLAS detector*, Eur. Phys. J. C **80** no. 2, (2020) 123, [arXiv:1908.08215 \[hep-ex\]](#).
- [61] A. Djouadi, S. Rosier-Lees, M. Bezouh, et al., *The Minimal Supersymmetric Standard Model: Group Summary Report*, 1999.
- [62] C. Abel, S. Afach, N. J. Ayres, et al., *Measurement of the Permanent Electric Dipole Moment of the Neutron*, Phys. Rev. Lett. **124** (2020) 081803.

- [63] P. Schmidt-Wellenburg, *The quest to find an electric dipole moment of the neutron*, [arXiv:1607.06609 \[hep-ex\]](#).
- [64] G. D'Ambrosio, G. Giudice, G. Isidori, and A. Strumia, *Minimal flavour violation: an effective field theory approach*, *Nuclear Physics B* **645** no. 1-2, (2002) 155–187.
- [65] E. Mobs, *The CERN accelerator complex in 2019*, 2019.
<https://cds.cern.ch/record/2684277>. General Photo.
- [66] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003.
- [67] CMS Collaboration, *The CMS experiment at the CERN LHC*, *JINST* **3** (2008) S08004.
- [68] LHCb Collaboration, *The LHCb Detector at the LHC*, *Journal of Instrumentation* **3** no. 08, (2008) S08005.
- [69] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, *Journal of Instrumentation* **3** no. 08, (2008) S08002.
- [70] A. Collaboration. *ATLAS luminosity public results Run 2*, 2019.
- [71] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report, 1*, 1997.
<https://cds.cern.ch/record/331063>.
- [72] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report, 2*, 1997.
<https://cds.cern.ch/record/331064>.
- [73] ATLAS Collaboration, *ATLAS Barrel Toroid: Technical Design Report*, 1997.
<https://cds.cern.ch/record/331065>.
- [74] ATLAS Collaboration, M. S. Alam et al., *ATLAS pixel detector: Technical design report*, Tech. Rep. CERN-LHCC-98-13, 5, 1998.
- [75] ATLAS Collaboration, M. Capeans, G. Darbo, K. Einsweiler, et al., *ATLAS Insertable B-Layer Technical Design Report*, tech. rep., 9, 2010.
- [76] ATLAS Collaboration, *Operation and performance of the ATLAS semiconductor tracker*, *JINST* **9** (2014) P08009, [arXiv:1404.7473 \[hep-ex\]](#).
- [77] B. Mindur, *ATLAS Transition Radiation Tracker (TRT): Straw tubes for tracking and particle identification at the Large Hadron Collider*, Tech. Rep.

- ATL-INDET-PROC-2016-001, CERN, Geneva, 2017.
<https://cds.cern.ch/record/2139567>.
- [78] ATLAS Collaboration, *ATLAS Calorimeter Performance: Technical Design Report*, 1996. <https://cds.cern.ch/record/331059>.
- [79] ATLAS Collaboration, *ATLAS Liquid-Argon Calorimeter: Technical Design Report*, 1996. <https://cds.cern.ch/record/331061>.
- [80] ATLAS Collaboration, *ATLAS Tile Calorimeter: Technical Design Report*, 1996. <https://cds.cern.ch/record/331062>.
- [81] A. Artamonov, D. Bailey, G. Belanger, et al., *The ATLAS Forward Calorimeter*, *Journal of Instrumentation* **3** no. 02, (2008) P02010.
- [82] ATLAS Collaboration, *ATLAS Muon Spectrometer: Technical Design Report*, 1997. <https://cds.cern.ch/record/331068>.
- [83] ATLAS Collaboration, *Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton–proton collision data*, *Eur. Phys. J. C* **74** (2014) 2941, [arXiv:1404.2240](https://arxiv.org/abs/1404.2240) [hep-ex].
- [84] ATLAS Collaboration, *Electron efficiency measurements with the ATLAS detector using 2012 LHC proton–proton collision data*, *Eur. Phys. J. C* **77** (2017) 195, [arXiv:1612.01456](https://arxiv.org/abs/1612.01456) [hep-ex].
- [85] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** no. 8, (2019) 639, [arXiv:1902.04655](https://arxiv.org/abs/1902.04655) [physics.ins-det].
- [86] ATLAS Collaboration, *Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton–proton collision data*, *Eur. Phys. J. C* **74** (2014) 3130, [arXiv:1407.3935](https://arxiv.org/abs/1407.3935) [hep-ex].
- [87] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **76** (2016) 292, [arXiv:1603.05598](https://arxiv.org/abs/1603.05598) [hep-ex].
- [88] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, *Eur. Phys. J. C* **77** (2017) 490, [arXiv:1603.02934](https://arxiv.org/abs/1603.02934) [hep-ex].

- [89] ATLAS Collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, *Eur. Phys. J. C* **77** (2017) 466, [arXiv:1703.10485 \[hep-ex\]](#).
- [90] M. Cacciari, G. P. Salam, and G. Soyez, *The anti-kt jet clustering algorithm*, *Journal of High Energy Physics* **2008** no. 04, (2008) 063–063.
- [91] ATLAS Collaboration, *Optimisation of large-radius jet reconstruction for the ATLAS detector in 13 TeV proton–proton collisions*, *Eur. Phys. J. C* **81** no. 4, (2021) 334, [arXiv:2009.04986 \[hep-ex\]](#).
- [92] ATLAS Collaboration, *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s}=8$ TeV using the ATLAS detector*, *Eur. Phys. J. C* **76** (2016) 581, 1510.
- [93] ATLAS Collaboration, *Reconstruction, Energy Calibration, and Identification of Hadronically Decaying Tau Leptons in the ATLAS Experiment for Run-2 of the LHC*, Tech. Rep. ATL-PHYS-PUB-2015-045, 2015. <https://cds.cern.ch/record/2064383>.
- [94] ATLAS Collaboration, *Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **78** (2018) 903, [arXiv:1802.08168 \[hep-ex\]](#).
- [95] C. Lester and D. Summers, *Measuring masses of semi-invisibly decaying particles pair produced at hadron colliders*, *Phys. Lett. B* **463** (1999) 99, [arXiv:9906349 \[hep-ph\]](#).
- [96] A. Barr, C. Lester, and P. Stephens, *$mT2$: the truth behind the glamour*, *J. Phys. G* **29** (2003), [arXiv:0304226 \[hep-ph\]](#).
- [97] ATLAS Collaboration, “Object-based missing transverse momentum significance in the ATLAS Detector.” ATLAS-CONF-2018-038, 2018. <https://cds.cern.ch/record/2630948>.
- [98] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*, *Eur. Phys. J. C* **70** (2010) 823, [arXiv:1005.4568 \[physics.ins-det\]](#).
- [99] T. Sjöstrand, S. Ask, J. R. Christiansen, et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159, [arXiv:1410.3012 \[hep-ph\]](#).
- [100] T. Gleisberg, S. Höche, F. Krauss, et al., *Event generation with SHERPA 1.1*,

- JHEP **02** (2009) 007, [arXiv:0811.4622 \[hep-ph\]](#).
- [101] S. Agostinelli et al., *GEANT4: A simulation toolkit*, Nucl. Instrum. Meth. A **506** (2003) 250.
- [102] E. Richter-Was, D. Froidevaux, and L. Poggioli, *ATLFAST 2.0 a fast simulation package for ATLAS*, Tech. Rep. ATL-PHYS-98-131, CERN, Geneva, 1998.
<https://cds.cern.ch/record/683751>.
- [103] W. Lukas, *Fast Simulation for ATLAS: Atlfast-II and ISF*, *Journal of Physics: Conference Series* **396** no. 2, (2012) 022031.
- [104] ATLAS Collaboration, “The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim.” ATL-PHYS-PUB-2010-013, 2010.
<https://cds.cern.ch/record/1300517>.
- [105] ROOT Collaboration, K. Cranmer, G. Lewis, L. Moneta, A. Shibata, and W. Verkerke, *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*, Tech. Rep. CERN-OPEN-2012-016, New York U., New York, 2012.
<https://cds.cern.ch/record/1456844>.
- [106] F. James and M. Roos, *Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations*, *Comput. Phys. Commun.* **10** (1975) 343–367.
- [107] P. Virtanen, R. Gommers, T. E. Oliphant, et al., *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, *Nature Methods* **17** (2020) 261–272.
- [108] Jerzy Neyman and Egon Sharpe Pearson, *IX. On the problem of the most efficient tests of statistical hypotheses*, *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* **231** (1933) 289–337.
- [109] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) 1554,
[arXiv:1007.1727 \[physics.data-an\]](#).
- [110] L. Lyons, *Statistical Issues in Searches for New Physics*, 2014.
- [111] A. L. Read, *Modified frequentist analysis of search results (the CL_s method)*,.

- [112] ATLAS Collaboration, *SimpleAnalysis: Truth-level Analysis Framework*, Tech. Rep. ATL-PHYS-PUB-2022-017, CERN, Geneva, 2022.
<https://cds.cern.ch/record/2805991>.
- [113] ATLAS Collaboration, *Reproducing searches for new physics with the ATLAS experiment through publication of full statistical likelihoods*, Tech. Rep. ATL-PHYS-PUB-2019-029, CERN, Geneva, 2019.
<https://cds.cern.ch/record/2684863>.
- [114] L. Heinrich, M. Feickert, and G. Stark, “PyHF: v0.7.0rc4.”
<https://doi.org/10.5281/zenodo.1169739>.
<https://github.com/scikit-hep/pyhf/releases/tag/v0.7.0rc4>.
- [115] L. Heinrich, M. Feickert, G. Stark, and K. Cranmer, *PyHF: pure-Python implementation of HistFactory statistical models*, *Journal of Open Source Software* **6** no. 58, (2021) 2823.
- [116] ATLAS Collaboration, *Implementation of simplified likelihoods in HistFactory for searches for supersymmetry*, Tech. Rep. ATL-PHYS-PUB-2021-038, CERN, Geneva, 2021. <https://cds.cern.ch/record/2782654>.
- [117] CMS Collaboration, *Simplified likelihood for the re-interpretation of public CMS results*, Tech. Rep. CMS-NOTE-2017-001, CERN, Geneva, 2017.
<https://cds.cern.ch/record/2242860>.
- [118] K. Cranmer and I. Yavin, *RECAST – extending the impact of existing analyses*, *Journal of High Energy Physics* **2011** (2011) 1029–8479.
- [119] ATLAS Collaboration, *RECAST framework reinterpretation of an ATLAS Dark Matter Search constraining a model of a dark Higgs boson decaying to two b -quarks*, Tech. Rep. ATL-PHYS-PUB-2019-032, CERN, Geneva, 2019.
<https://cds.cern.ch/record/2686290>.
- [120] D. Merkel, *Docker: lightweight linux containers for consistent development and deployment*, *Linux journal* **2014** no. 239, (2014) 2.
- [121] ATLAS Collaboration, *New techniques for jet calibration with the ATLAS detector*, [arXiv:2303.17312](https://arxiv.org/abs/2303.17312) [hep-ex].
- [122] ATLAS Collaboration, *Constituent-Based Top-Quark Tagging with the ATLAS Detector*, Tech. Rep. ATL-PHYS-PUB-2022-039, CERN, Geneva, 2022.

- <https://cds.cern.ch/record/2825328>.
- [123] ATLAS Collaboration, *Search for direct pair production of sleptons and charginos decaying to two leptons and neutralinos with mass splittings near the W -boson mass in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, [arXiv:2209.13935 \[hep-ex\]](#).
- [124] K. Hornik, M. Stinchcombe, and H. White, *Multilayer feedforward networks are universal approximators*, [Neural Networks](#) **2** no. 5, (1989) 359–366.
- [125] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Learning representations by back-propagating errors*, [Nature](#) **323** (1986) 533–536.
- [126] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980 \[cs.LG\]](#).
- [127] P. J. Huber [Ann. Math. Statist.](#) **35** (1964) 73.
- [128] ATLAS Collaboration, *Summary of the ATLAS experiment’s sensitivity to supersymmetry after LHC Run 1 – interpreted in the phenomenological MSSM.*, [J. High Energ. Phys.](#) **2015** (2015) 134.
- [129] ATLAS Collaboration, *Dark matter interpretations of ATLAS searches for the electroweak production of supersymmetric particles in $\sqrt{s} = 8$ TeV proton-proton collisions.*, [J. High Energ. Phys.](#) **2016** (2016) 175.
- [130] CMS Collaboration, V. Khachatryan et al., *Phenomenological MSSM interpretation of CMS searches in pp collisions at $\sqrt{s} = 7$ and 8 TeV*, [JHEP](#) **10** (2016) 129, [arXiv:1606.03577 \[hep-ex\]](#).
- [131] GAMBIT Collaboration, P. Athron et al., *GAMBIT: The Global and Modular Beyond-the-Standard-Model Inference Tool*, [Eur. Phys. J. C](#) **77** no. 11, (2017) 784, [arXiv:1705.07908 \[hep-ph\]](#).
- [132] GAMBIT Collaboration, P. Athron et al., *A global fit of the MSSM with GAMBIT*, [Eur. Phys. J. C](#) **77** no. 12, (2017) 879, [arXiv:1705.07917 \[hep-ph\]](#).
- [133] GAMBIT Collaboration, V. Ananyev et al., *Collider constraints on electroweakinos in the presence of a light gravitino*, [Eur. Phys. J. C](#) **83** no. 6, (2023) 493, [arXiv:2303.09082 \[hep-ph\]](#).
- [134] E. Bagnaschi et al., *Likelihood Analysis of the p MSSM11 in Light of LHC 13-TeV*

- Data*, *Eur. Phys. J. C* **78** no. 3, (2018) 256, [arXiv:1710.11091 \[hep-ph\]](#).
- [135] S. Kraml, S. Kulkarni, U. Laa, et al., *SModelS: a tool for interpreting simplified-model results from the LHC and its application to supersymmetry*, *Eur. Phys. J. C* **74** (2014) 2868, [arXiv:1312.4175 \[hep-ph\]](#).
- [136] F. Ambrogio et al., *SModelS v1.2: long-lived particles, combination of signal regions, and other novelties*, *Comput. Phys. Commun.* **251** (2020) 106848, [arXiv:1811.10624 \[hep-ph\]](#).
- [137] M. Van Beekveld, W. Beenakker, M. Schutten, and J. De Wit, *Dark matter, fine-tuning and $(g - 2)_\mu$ in the pMSSM*, *SciPost Phys.* **11** no. 3, (2021) 049, [arXiv:2104.03245 \[hep-ph\]](#).
- [138] LEPSUSYWG, ALEPH, DELPHI, L3 and OPAL experiments, “Combined lep chargino results.” <https://lepsusy.web.cern.ch/lepsusy/>.
- [139] ATLAS Collaboration, *Search for charginos and neutralinos in final states with two boosted hadronically decaying bosons and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **104** no. 11, (2021) 112010, [arXiv:2108.07586 \[hep-ex\]](#).
- [140] ATLAS Collaboration, *Search for direct production of electroweakinos in final states with one lepton, missing transverse momentum and a Higgs boson decaying into two b-jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, [arXiv:1909.09226 \[hep-ex\]](#).
- [141] ATLAS Collaboration, *Search for electroweak production of charginos and sleptons decaying into final states with two leptons and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions using the ATLAS detector*, *Eur. Phys. J. C* **80** (2020) 123, [arXiv:1908.08215 \[hep-ex\]](#).
- [142] ATLAS Collaboration, *Searches for new phenomena in events with two leptons, jets, and missing transverse momentum in 139 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, [arXiv:2204.13072 \[hep-ex\]](#).
- [143] ATLAS Collaboration, *Searches for electroweak production of supersymmetric particles with compressed mass spectra in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Phys. Rev. D* **101** (2020) 052005, [arXiv:1911.12606 \[hep-ex\]](#).
- [144] ATLAS Collaboration, *Search for chargino–neutralino pair production in final*

- states with three leptons and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Eur. Phys. J. C* **81** (2021) 1118.
- [145] ATLAS Collaboration, *Search for supersymmetry in events with four or more charged leptons in 139 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *J. High Energ. Phys.* **2021** (2021) 167.
- [146] ATLAS Collaboration, *Search for long-lived charginos based on a disappearing-track signature using 136 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **82** no. 7, (2022) 606, [arXiv:2201.02472 \[hep-ex\]](#).
- [147] ATLAS Collaboration, *Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, *Phys. Rev. D* **101** no. 1, (2020) 012002, [arXiv:1909.02845 \[hep-ex\]](#).
- [148] ATLAS Collaboration, *Combination of searches for invisible decays of the Higgs boson using 139 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, *Phys. Lett. B* **842** (2023) 137963, [arXiv:2301.10731 \[hep-ex\]](#).
- [149] J. Dickinson, S. Bein, S. Heinemeyer, et al., *A Grand Scan of the pMSSM Parameter Space for Snowmass 2021*, 2022.
- [150] W. Porod, *SPheno, a program for calculating supersymmetric spectra, SUSY particle decays and SUSY particle production at $e+e^-$ colliders*, *Computer Physics Communications* **153** no. 2, (2003) 275–315.
- [151] W. Porod and F. Staub, *SPheno 3.1: extensions including flavour, CP-phases and models beyond the MSSM*, *Computer Physics Communications* **183** no. 11, (2012) 2458–2469.
- [152] H. Bahl, T. Hahn, S. Heinemeyer, et al., *Precision calculations in the MSSM Higgs-boson sector with FeynHiggs 2.14*, 2018. <https://arxiv.org/abs/1811.09073>.
- [153] H. Bahl, S. Heinemeyer, W. Hollik, and G. Weiglein, *Reconciling EFT and hybrid calculations of the light MSSM Higgs-boson mass*, *The European Physical Journal C* **78** no. 1, (2018).

- [154] H. Bahl and W. Hollik, *Precise prediction for the light MSSM Higgs-boson mass combining effective field theory and fixed-order calculations*, *The European Physical Journal C* **76** no. 9, (2016).
- [155] T. Hahn, S. Heinemeyer, W. Hollik, H. Rzehak, and G. Weiglein, *High-Precision Predictions for the light CP-Even Higgs Boson Mass of the Minimal Supersymmetric Standard Model*, *Physical Review Letters* **112** no. 14, (2014).
- [156] M. Frank, T. Hahn, S. Heinemeyer, et al., *The Higgs boson masses and mixings of the complex MSSM in the Feynman-diagrammatic approach*, *Journal of High Energy Physics* **2007** no. 02, (2007) 047–047.
- [157] G. Degrandi, S. Heinemeyer, W. Hollik, P. Slavich, and G. Weiglein, *Towards high-precision predictions for the MSSM Higgs sector*, *The European Physical Journal C* **28** no. 1, (2003) 133–143.
- [158] S. Heinemeyer, W. Hollik, and G. Weiglein, *The masses of the neutral CP-even Higgs bosons in the MSSM: Accurate analysis at the two-loop level*, *The European Physical Journal C* **9** no. 2, (1999) 343–366.
- [159] S. Heinemeyer, W. Hollik, and G. Weiglein, *FeynHiggs: a program for the calculation of the masses of the neutral -even Higgs bosons in the MSSM*, *Computer Physics Communications* **124** no. 1, (2000) 76–89.
- [160] D. Barducci, G. Belanger, J. Bernon, et al., *Collider limits on new physics within micrOMEGAs4.3*, 2016. <https://arxiv.org/abs/1606.03834>.
- [161] G. Belanger, A. Mjallal, and A. Pukhov, *Recasting direct detection limits within micrOMEGAs and implication for non-standard dark matter scenarios*, *The European Physical Journal C* **81** no. 3, (2021).
- [162] W. Beenakker, R. Hoepker, and M. Spira, *PROSPINO: A Program for the Production of Supersymmetric Particles in Next-to-leading Order QCD*, 1996. <https://arxiv.org/abs/hep-ph/9611232>.
- [163] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, *MadGraph 5 : Going Beyond*, *JHEP* **06** (2011) 128, [arXiv:1106.0522](https://arxiv.org/abs/1106.0522) [hep-ph].
- [164] J. Alwall, R. Frederix, S. Frixione, et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079, [arXiv:1405.0301](https://arxiv.org/abs/1405.0301) [hep-ph].

- [165] E. Maguire, L. Heinrich, and G. Watt, *HEPData: a repository for high energy physics data*, *J. Phys. Conf. Ser.* **898** no. 10, (2017) 102006, [arXiv:1704.05473 \[hep-ex\]](#).
- [166] K. Sinha, *Non-thermal dark matter: Prospects and challenges*, *AIP Conf. Proc.* **1604** no. 1, (2015) 105–114, [arXiv:1311.0884 \[hep-ph\]](#).
- [167] A. Buckley, *PySLHA: a Pythonic interface to SUSY Les Houches Accord data*, *Eur. Phys. J. C* **75** no. 10, (2015) 467, [arXiv:1305.4194 \[hep-ph\]](#).
- [168] ATLAS Collaboration, *Search for direct production of electroweakinos in final states with one lepton, missing transverse momentum and a Higgs boson decaying into two b-jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **80** (2020) 691.
- [169] D. R. Tovey, *On measuring the masses of pair-produced semi-invisibly decaying particles at hadron colliders*, *Journal of High Energy Physics* **2008** no. 04, (2008) 034–034.
- [170] G. Polesello and D. R. Tovey, *Supersymmetric particle mass measurement with the boost-corrected contranverse mass*, *Journal of High Energy Physics* **2010** no. 3, (2010).
- [171] ATLAS Collaboration, *Searches for electroweak production of supersymmetric particles with compressed mass spectra in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Phys. Rev. D* **101** (2020) 052005.
- [172] A. Djouadi, M. M. Muhlleitner, and M. Spira, *Decays of supersymmetric particles: The Program SUSY-HIT (SUSpect-SdecaY-Hdecay-InTerface)*, *Acta Phys. Polon. B* **38** (2007) 635–644, [arXiv:hep-ph/0609292](#).
- [173] ATLAS Collaboration, *Measurement of the Higgs boson mass in the $H \rightarrow ZZ^* \rightarrow 4\ell$ decay channel using 139 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions recorded by the ATLAS detector at the LHC*, [arXiv:2207.00320 \[hep-ex\]](#).
- [174] LHCb Collaboration, R. Aaij et al., *Measurement of the $B_s^0 \rightarrow \mu^+ \mu^-$ decay properties and search for the $B^0 \rightarrow \mu^+ \mu^-$ and $B_s^0 \rightarrow \mu^+ \mu^- \gamma$ decays*, *Phys. Rev. D* **105** no. 1, (2022) 012010, [arXiv:2108.09283 \[hep-ex\]](#).
- [175] J. Haller, A. Hoecker, R. Kogler, et al., *Update of the global electroweak fit and constraints on two-Higgs-doublet models*, *Eur. Phys. J. C* **78** no. 8, (2018) 675,

- [arXiv:1803.01853 \[hep-ph\]](#).
- [176] ALEPH, DELPHI, L3, OPAL, SLD, LEP Electroweak Working Group, SLD Electroweak Group, SLD Heavy Flavour Group Collaboration, S. Schael et al., *Precision electroweak measurements on the Z resonance*, **Phys. Rept.** **427** (2006) 257–454, [arXiv:hep-ex/0509008](#).
- [177] CDF Collaboration, T. Aaltonen et al., *High-precision measurement of the W boson mass with the CDF II detector*, **Science** **376** no. 6589, (2022) 170–176.
- [178] S. Heinemeyer, W. Hollik, G. Weiglein, and L. Zeune, *Implications of LHC search results on the W boson mass prediction in the MSSM*, **JHEP** **12** (2013) 084, [arXiv:1311.1663 \[hep-ph\]](#).
- [179] M. E. Peskin and T. Takeuchi, *New constraint on a strongly interacting Higgs sector*, **Phys. Rev. Lett.** **65** (1990) 964–967.
- [180] M. E. Peskin and T. Takeuchi, *Estimation of oblique electroweak corrections*, **Phys. Rev. D** **46** (1992) 381–409.
- [181] LZ Collaboration, J. Aalbers et al., *First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment*, [arXiv:2207.03764 \[hep-ex\]](#).
- [182] PICO Collaboration, C. Amole et al., *Dark Matter Search Results from the Complete Exposure of the PICO-60 C₃F₈ Bubble Chamber*, **Phys. Rev. D** **100** no. 2, (2019) 022001, [arXiv:1902.04031 \[astro-ph.CO\]](#).
- [183] ATLAS Collaboration, *Search for direct production of charginos, neutralinos and sleptons in final states with two leptons and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, **JHEP** **05** (2014) 071, [arXiv:1403.5294 \[hep-ex\]](#).
- [184] ATLAS Collaboration, *Search for direct production of electroweakinos in final states with one lepton, jets and missing transverse momentum and in p p collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2022-059, CERN, Geneva, 2022. <http://cds.cern.ch/record/2826702>.
- [185] J. Bradbury, R. Frostig, P. Hawkins, et al., *JAX: composable transformations of Python+NumPy programs*, 2018. <http://github.com/google/jax>.
- [186] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*, **JINST** **14**

- no. 12, (2019) P12006, [arXiv:1908.00005 \[hep-ex\]](#).
- [187] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** no. 7, (2021) 578, [arXiv:2012.00578 \[hep-ex\]](#).
- [188] M. Baak, G. J. Besjes, D. Côte, et al., *HistFitter software framework for statistical data analysis*, *Eur. Phys. J. C* **75** (2015) 153, [arXiv:1410.1280 \[hep-ex\]](#).
- [189] T. Sjostrand, S. Mrenna, and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852–867, [arXiv:0710.3820 \[hep-ph\]](#).
- [190] R. D. Ball et al., *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244, [arXiv:1207.1303 \[hep-ph\]](#).
- [191] ATLAS Collaboration, “The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie–Landshoff diffractive model.” ATL-PHYS-PUB-2016-017, 2016.
<https://cds.cern.ch/record/2206965>.
- [192] ATLAS Collaboration, “Improvements in $t\bar{t}$ modelling using NLO+PS Monte Carlo generators for Run 2.” ATL-PHYS-PUB-2018-009, 2018.
<https://cds.cern.ch/record/2630327>.
- [193] ATLAS Collaboration, “Multi-Boson Simulation for 13 TeV ATLAS Analyses.” ATL-PHYS-PUB-2017-005, 2017. <https://cds.cern.ch/record/2261933>.
- [194] ATLAS Collaboration, “ATLAS simulation of boson plus jets processes in Run 2.” ATL-PHYS-PUB-2017-006, 2017. <https://cds.cern.ch/record/2261937>.
- [195] P. Nason, *A new method for combining NLO QCD with shower Monte Carlo algorithms*, *JHEP* **11** (2004) 040, [arXiv:hep-ph/0409146](#).
- [196] S. Frixione, P. Nason, and C. Oleari, *Matching NLO QCD computations with parton shower simulations: the POWHEG method*, *JHEP* **11** (2007) 070, [arXiv:0709.2092 \[hep-ph\]](#).
- [197] S. Alioli, P. Nason, C. Oleari, and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) 043, [arXiv:1002.2581 \[hep-ph\]](#).
- [198] M. Czakon and A. Mitov, *Top++: A program for the calculation of the top-pair*

- cross-section at hadron colliders*, *Comput. Phys. Commun.* **185** (2014) 2930, [arXiv:1112.5675 \[hep-ph\]](#).
- [199] NNPDF Collaboration, R. D. Ball et al., *Parton distributions for the LHC Run II*, *JHEP* **04** (2015) 040, [arXiv:1410.8849 \[hep-ph\]](#).
- [200] ATLAS Collaboration, *Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **09** (2014) 145, [arXiv:1406.3660 \[hep-ex\]](#).
- [201] H.-L. Lai et al., *New parton distributions for collider physics*, *Phys. Rev. D* **82** (2010) 074024, [arXiv:1007.2241 \[hep-ph\]](#).
- [202] J. Pumplin et al., *New Generation of Parton Distributions with Uncertainties from Global QCD Analysis*, *JHEP* **07** (2002) 012, [arXiv:hep-ph/0201195](#).
- [203] S. Höche, F. Krauss, S. Schumann, and F. Siegert, *QCD matrix elements and truncated showers*, *JHEP* **05** (2009) 053, [arXiv:0903.1219 \[hep-ph\]](#).
- [204] E. Bothmann et al., *Event generation with Sherpa 2.2*, *SciPost Phys.* **7** no. 3, (2019) 034, [arXiv:1905.09127 \[hep-ph\]](#).
- [205] ATLAS Collaboration, “Monte Carlo Generators for the Production of a W or Z/γ^* Boson in Association with Jets at ATLAS in Run 2.” ATL-PHYS-PUB-2016-003, 2016. <https://cds.cern.ch/record/2120133>.
- [206] R. Gavin, Y. Li, F. Petriello, and S. Quackenbush, *FEWZ 2.0: A code for hadronic Z production at next-to-next-to-leading order*, [arXiv:1011.3540 \[hep-ph\]](#).
- [207] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data*, *JINST* **14** (2019) P12006, [arXiv:1908.00005 \[hep-ex\]](#).
- [208] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 578, [arXiv:2012.00578 \[hep-ex\]](#).
- [209] ATLAS Collaboration, *ATLAS b -jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 970, [arXiv:1907.05120 \[hep-ex\]](#).
- [210] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, *Self-Normalizing*

- Neural Networks*, CoRR (2017), [arXiv:1706.02515](#).
- [211] J. L. Ba, J. R. Kiros, and G. E. Hinton, *Layer Normalization*, [arXiv:1607.06450](#) [stat.ML].
- [212] F. Chollet et al., “Keras.” <https://keras.io>, 2015.
- [213] M. Abadi, A. Agarwal, P. Barham, et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, 2015. <http://tensorflow.org/>. Software available from tensorflow.org.
- [214] K. He, X. Zhang, S. Ren, and J. Sun, *Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification*, [arXiv:1502.01852](#) [cs.CV].
- [215] ATLAS Collaboration, “Formulae for Estimating Significance.” ATL-PHYS-PUB-2020-025, 2020. <https://cds.cern.ch/record/2736148>.
- [216] M. Cuturi, *Sinkhorn Distances: Lightspeed Computation of Optimal Transportation Distances*, [arXiv:1306.0895](#) [stat.ML].
- [217] D. Nix and A. Weigend, *Estimating the mean and variance of the target probability distribution*, in *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, pp. , 55–60 vol.1. 1994.
- [218] C. Bishop, *Mixture density networks*, tech. rep., Aston University, 1994.
- [219] A. Buckley, “The hepthesis L^AT_EX class.”.

List of figures

2.1. Fermion and scalar loop contributions to the Higgs mass.	12
2.2. The observed and expected exclusion limits reported by an ATLAS search for pair-production of charginos with a leptonic final state [60], interpreted in a simplified model.	19
3.1. The particle accelerator complex at CERN [65].	25
3.2. Displayed are: (a) Total integrated luminosity delivered by the LHC (green), recorded by ATLAS (yellow) and good for physics analysis (blue) during Run 2; (b) the luminosity weighted distribution of the mean number of interactions per bunch crossing for each year of data-taking during Run 2. Reproduced from Reference [70].	26
3.3. Computer-generated image of the ATLAS detector, including each sub-detector. Image from Reference [66].	27
3.4. Schematic diagram of the ATLAS inner detector. Image from Reference [66].	29
3.5. Schematic diagram of the ATLAS calorimeter system. Image from Reference [66].	30
3.6. Schematic diagram of the ATLAS muon spectrometer. Image from Reference [66].	31
3.7. Observed and expected exclusion limits on a SUSY simplified model of chargino-pair production, reproduced from Reference [60].	48
3.8. The basic neural network structure of an input layer, hidden layers and an output layer.	51

- 4.1. Workflow for the pMSSM scan, from sampling the pMSSM parameter space to determining if each model is excluded. 59
- 4.2. Comparison of detector-level and truth-level $CL_s^{\text{Exp.}}$ for a test sample of pMSSM models. The points are coloured by the analysis result being displayed, which is the analysis with the lowest $CL_s^{\text{Exp.}}$ for that model. A vertical line at 0.05 on the x -axis indicates the detector-level exclusion boundary. Horizontal lines at 0.001 and 0.1 on the y -axis indicate the final decision boundary for evaluating a model at detector-level. The red regions indicate CL_s scenarios which would result in the false categorisation of a model. 65
- 4.3. Scatter of models from the EWKino scan in the $m(\tilde{\chi}_1^0)$ vs. Ωh^2 plane. (a) is coloured with RGB value set to the wino/higgsino/bino fraction of the LSP and (b) is coloured by the dominant annihilation mechanism. The experimentally measured value of the relic density [33], $\Omega h^2 = 0.12$, is indicated by a horizontal dashed line. 72
- 4.4. Scatter plot of models from the Bino-DM scan in the $m(\tilde{\chi}_1^0)$ vs. Ωh^2 plane. (a) is coloured by the dominant annihilation mechanism and (b) is coloured by the LSP higgsino fraction with black representing 0. The experimentally measured value of the relic density [33], $\Omega h^2 = 0.12$, is indicated by a horizontal dashed line. 73
- 4.5. Mass-spectrum for two Bino-DM scan models: (a) A model in the funnel region with a bino/higgsino mixed LSP and higgsino-like $\tilde{\chi}_1^\pm$, (b) A model in the bulk region with an almost pure-bino LSP and pure-wino $\tilde{\chi}_1^\pm$. Produced using PySLHA [167]. 76
- 4.6. Diagrams of the signatures targeted by the fully hadronic (FullHad) analysis. 77

- 4.7. (a) Fraction of models excluded by FullHad in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane for the EWKino scan. Grey bins contain no models and black bins have 0% exclusion. The observed exclusion contour on a simplified model of pure-wino $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W \tilde{\chi}_1^0 Z \tilde{\chi}_1^0$ and $pp \rightarrow \tilde{\chi}_1^\pm \tilde{\chi}_1^\pm \rightarrow Z \tilde{\chi}_1^0 Z \tilde{\chi}_1^0$ with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z \tilde{\chi}_1^0) = 100\%$ and $\text{BR}(\tilde{\chi}_1^\pm \rightarrow W \tilde{\chi}_1^0) = 100\%$ is overlaid.
 (b) The distribution of the $\tilde{\chi}_1^\pm$ wino fraction for models which lie inside the FullHad simplified model exclusion contour, shown before and after applying FullHad. The ratio panel shows the fraction of models excluded in each bin. 78
- 4.8. EWKino scan models in the plane of lightest chargino mass vs. the cross-section for pair-production of the lightest chargino. The points are coloured by the wino fraction of the lightest chargino. 79
- 4.9. (a) Fraction of EWKino scan models excluded by FullHad in the $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) - m(\tilde{\chi}_1^0)$ plane. Grey bins contain no models and black bins have 0% exclusion.
 (b) Total inclusive truth-smeared FullHad analysis signal region yield for each SUSY production process separated by signal region, for all models in the EWKino scan which are excluded at truth-level by FullHad and satisfy $m(\tilde{\chi}_1^\pm) < 400$ GeV, $m(\tilde{\chi}_1^0) < 400$ GeV and $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) < 50$ GeV. 79
- 4.10. Diagram of the signature targeted by the 1Lbb analysis. 81
- 4.11. (a) Fraction of EWKino scan models excluded by 1Lbb in the $\tilde{\chi}_2^0 - \tilde{\chi}_1^0$ mass plane, overlain with the observed exclusion contour for the simplified model in Figure 4.10. Grey bins contain no models and black bins have 0% exclusion.
 (b) Distribution of $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h)$ before and after 1Lbb for EWKino models which lie inside the observed 1Lbb exclusion contour in the $\tilde{\chi}_2^0 - \tilde{\chi}_1^0$ mass plane. The ratio plot shows the fraction of models excluded by 1Lbb in each bin. 81
- 4.12. Diagrams of the signatures targeted by the 2L0J analysis. 82
- 4.13. Fraction of models from the EWKino scan excluded by 2L0J in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane. Grey bins contain no models and black bins have 0% exclusion. The observed exclusion contour for the simplified model in Figure 4.12a is overlain. 84

- 4.14. Two plots in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane: (a) shows each Bino-DM scan model coloured by the $\tilde{\chi}_1^\pm$ higgsino fraction, (b) shows the fraction of Bino-DM models excluded by 2L0J with the 2L0J simplified model exclusion contour overlain. Grey bins contain no models and black bins have 0% exclusion. 84
- 4.15. Diagrams of the signatures targeted by the compressed analysis. 85
- 4.16. Fraction of EWKino scan models with a higgsino-like LSP that are excluded by the Compressed analysis, shown in the $m(\tilde{\chi}_2^0)$ vs. $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane. Grey bins contain no models and black bins have 0% exclusion. The observed higgsino simplified model exclusion contour is overlain. 86
- 4.17. Fraction of Bino-DM scan models that are excluded by the Compressed analysis, shown in the $m(\tilde{\chi}_2^0)$ vs. $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane. (a) shows models with $M_1 \times M_2 > 0$ and (b) shows models with $M_1 \times M_2 < 0$. Grey bins contain no models and black bins have 0% exclusion. The plots are overlain with the relevant observed simplified model exclusion contour for each case. . . 86
- 4.18. Scatter plot of the mass difference between the second and lightest neutralino versus (a) $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 \mu^+ \mu^-)$, (b) $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 + \gamma)$. The plots include pMSSM models from both the EWKino and Bino-DM scans that lie inside the relevant higgsino or wino/bino simplified model expected exclusion contour from the Compressed analysis. Points are coloured with RGB value set to the $\tilde{\chi}_2^0$ wino/higgsino/bino fraction. The crosses in (a) indicate the higgsino-LSP simplified model branching ratio from SUSYHIT (green), the branching ratio used in the Compressed analysis for higgsino-LSP simplified models (grey), and the branching ratio used in the Compressed analysis for wino-LSP simplified models (red). 88
- 4.19. Diagram of the signature targeted by the 2L2J analysis. 89
- 4.20. Fraction of EWKino scan models excluded by 2L2J in several mass planes. Grey bins contain no models and black bins have 0% exclusion. (a) is overlain with the observed exclusion contour for the simplified model in Figure 4.19. 91

4.21. (a) Total inclusive truth-smeared 2L2J analysis signal region yield for each SUSY production process separated by signal region, for all models which are excluded by 2L2J (at truth-level) and lie outside the simplified model exclusion contour in the $\tilde{\chi}_2^0\text{--}\tilde{\chi}_1^0$ mass plane with $m(\tilde{\chi}_1^\pm) < 300$ GeV. (b) Distribution of the $\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0$ branching ratio for models in the EWKino scan which lie inside the 2L2J expected exclusion contour. The ratio panel shows the fraction of models excluded by 2L2J.	92
4.22. Diagrams of the signatures targeted by the 3L analysis.	93
4.23. Fraction of models excluded by the 3L onshell WZ signal regions in (a) the $\tilde{\chi}_1^\pm\text{--}\tilde{\chi}_1^0$ mass plane for bino-LSP models from the EWKino scan and (b) the $m(\tilde{\chi}_1^\pm) - \Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane for higgsino-LSP models from the EWKino scan. Grey bins contain no models and black bins have 0% exclusion. The observed exclusion contour for the relevant simplified model is overlain. .	94
4.24. Distribution of the $\tilde{\chi}_2^0 \rightarrow Z\tilde{\chi}_1^0$ branching ratio for EWKino scan models which lie inside the 3L onshell (WZ) simplified model exclusion contour. The ratio plot shows the fraction of models excluded by 3L onshell (WZ).	94
4.25. Fraction of models excluded by the 3L (offshell) analysis in the 2D plane of $m(\tilde{\chi}_2^0)$ vs. $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$. Grey bins contain no models and black bins have 0% exclusion. The observed exclusion contour is overlain for the relevant simplified model in Table 4.11 for (a) bino-LSP models from the Bino-DM scan and (b) higgsino-LSP models from the EWKino scan.	95
4.26. Diagrams of the signatures targeted by the 4L analysis.	96
4.27. (a) Fraction of EWKino scan models excluded by 4L in the $\tilde{\chi}_1^\pm\text{--}\tilde{\chi}_1^0$ mass plane. Grey bins contain no models and black bins have 0% exclusion. (b) Total inclusive truth-smeared 4L analysis signal region yield for each SUSY production process, summed over all models in the EWKino scan that 4L is sensitive to (at truth-level).	96
4.28. Diagram of the signature targeted by the Disappearing Track analysis. .	97

- 4.29. (a) Scatter of pMSSM models in the chargino mass–lifetime plane. Both the EWKino and Bino-DM scans are included. Models coloured in blue are excluded by the disappearing track cross-section upper limits, in green are not excluded and in orange are outside the valid interpolation region. (b) Shows the fraction of EWKino and Bino-DM scan models excluded by the disappearing track analysis in the $m(\tilde{\chi}_1^\pm) - \Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ plane. Black bins have 0% exclusion. 98
- 4.30. Distribution of $\text{BR}(h \rightarrow \tilde{\chi}_1^0 \tilde{\chi}_1^0)$ (first row) and $m(A)$ (second row) before and after the ATLAS Run 2 direct searches. The first column shows the EWKino scan and the second column the Bino-DM scan. Relevant upper/lower limits are shown as red lines. Only models with $m(\tilde{\chi}_1^\pm) < 1200$ GeV and $m(\tilde{\chi}_1^0) < 500$ GeV are included to focus on the region the ATLAS searches are sensitive to. 100
- 4.31. Percentage overlaps in the exclusion of each ATLAS analysis. For each square, the number and colour indicates the percentage of the models excluded by the analysis on the y -axis which are also excluded by the analysis on the x -axis, ie. $100 \times N_{xy}/N_y$ where N_{xy} is the number of models excluded by the results on the x AND y axes and N_y is the number of models excluded by the result on the y axis. 101
- 4.32. Scatter plots coloured by the search (not including $m(A)$ and $\text{BR}(h \rightarrow \text{inv.})$ constraints) with unique (expected) sensitivity. Models that multiple searches are sensitive to are shown in silver. Only models that are excluded are shown and for clarity in (c) only models with unique analysis sensitivity are shown due to the density of points. The first row shows the EWKino scan and the second row the Bino-DM scan. The first column shows the $m(\tilde{\chi}_1^\pm) - m(\tilde{\chi}_1^0)$ plane and the second column the $m(\tilde{\chi}_2^0) - \Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane. 102
- 4.33. Distribution of the Higgs boson mass before and after ATLAS Run 2 for the EWKino scan model set. The red dashed lines indicate the 120 GeV–130 GeV mass window used to constrain the model set. The ratio panel shows the fraction of models excluded by ATLAS. Only models with $m(\tilde{\chi}_1^\pm) < 1200$ GeV and $m(\tilde{\chi}_1^0) < 500$ GeV are shown to focus on the region ATLAS searches are sensitive to. 104

- 4.34. Distribution of flavour observables for EWKino scan models before and after ATLAS Run 2. The red dashed lines indicate the window used when applying each constraint to the model sets. The ratio panel shows the fraction of models excluded by ATLAS. Only models with $m(\tilde{\chi}_1^\pm) < 1200$ GeV and $m(\tilde{\chi}_1^0) < 500$ GeV are shown to focus on the region ATLAS searches are sensitive to. 106
- 4.35. Distribution of electroweak precision observables for EWKino scan models before and after ATLAS Run 2. The red dashed lines indicate the bound(s) used when applying each constraint to the model sets. The ratio panel shows the fraction of models excluded by ATLAS. Only models with $m(\tilde{\chi}_1^\pm) < 1200$ GeV and $m(\tilde{\chi}_1^0) < 500$ GeV are shown to focus on the region ATLAS searches are sensitive to. 107
- 4.36. Percentage overlaps in the exclusion of each constraint. For each square, the number and colour indicate the percentage of the models excluded by the result on the y -axis which are also excluded by the result on the x -axis, ie. $100 \times N_{xy}/N_y$ where N_{xy} is the number of models excluded by the results on the x AND y axes and N_y is the number of models excluded by the result on the y axis. 108
- 4.37. Fraction of models excluded by Run 2 ATLAS results in the $\tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ mass plane for the (a) EWKino and (b) Bino-DM scans. Only models which pass all non-dark-matter external constraints are included. Grey bins contain no models and black bins have 0% exclusion. The overlain contour shows the envelope of the FullHad + 2L2J + 3L exclusion of a $\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W \tilde{\chi}_1^0 Z \tilde{\chi}_1^0$ simplified model with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^{(*)} \tilde{\chi}_1^0) = 100\%$ 110
- 4.38. Fraction of models excluded by Run 2 ATLAS results in the $\tilde{\chi}_2^0 - \tilde{\chi}_1^0$ mass plane for the (a) EWKino and (b) Bino-DM scans. Only models which pass all non-dark-matter external constraints are included. Grey bins contain no models and black bins have 0% exclusion. The overlain contour shows the envelope of the FullHad + 2L2J + 3L exclusion of a $\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W \tilde{\chi}_1^0 Z \tilde{\chi}_1^0$ simplified model with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^{(*)} \tilde{\chi}_1^0) = 100\%$ 110

- 4.39. Fraction of models excluded by Run 2 ATLAS results in the $m(\tilde{\chi}_2^0)$ vs. $\Delta m(\tilde{\chi}_2^0, \tilde{\chi}_1^0)$ plane for the (a) EWKino and (b) Bino-DM scans. Only models which pass all non-dark-matter external constraints are included. Grey bins contain no models and black bins have 0% exclusion. The overlain contours show the ATLAS 3L + Compressed exclusion of a $\tilde{\chi}_1^\pm \tilde{\chi}_2^0 \rightarrow W^* \tilde{\chi}_1^0 Z^* \tilde{\chi}_1^0$ simplified model with $\text{BR}(\tilde{\chi}_2^0 \rightarrow Z^{(*)} \tilde{\chi}_1^0) = 100\%$ and (a) a pure higgsino LSP and $\tilde{\chi}_2^0$, (b) a pure bino LSP and pure wino $\tilde{\chi}_2^0$ 112
- 4.40. Fraction of models excluded by Run 2 ATLAS results in the $m(\tilde{\chi}_1^\pm)$ vs. $\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$ plane for the (a) EWKino and (b) Bino-DM scans. Only models which pass all non-dark-matter external constraints are included. Grey bins contain no models and black bins have 0% exclusion. 112
- 4.41. Distribution of the LSP mass (first row) and $\tilde{\chi}_1^\pm$ mass (second row) before and after ATLAS Run 2, split by bino-/wino-/higgsino-type (note that in (b) all LSPs are bino-like by construction). The first column shows models in the EWKino scan that pass all non-dark-matter constraints and the second column models in the Bino-DM scan that pass all non-dark-matter constraints. The ratio panels show the fraction of models excluded by ATLAS in each case. 114
- 4.42. Summary plots showing the fraction of models excluded by ATLAS across the masses of each electroweakino for models which pass non-dark-matter external constraints. (a) shows the EWKino scan and (b) the Bino-DM scan. Grey bins contain no models and black bins have 0% exclusion. . . 115
- 4.43. $\tilde{\chi}_2^0$ branching ratio distributions for the EWKino scan (first column) and Bino-DM scan (second column) before and after ATLAS Run 2. The EWKino scan distributions are split by bino/wino LSP type. Only models that pass all non-dark-matter constraints are included. The ratio panels show the fraction of models excluded by ATLAS. Note that $\text{BR}(\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^\pm W^\mp)$ refers to the sum of the W^+ and W^- modes. 116

- 4.44. Models from the Bino-DM scan which pass all non-dark-matter external constraints, coloured by the dominant annihilation mechanism, before (first column) and after (second column) ATLAS Run 2. (a) and (b) show the $m(\tilde{\chi}_1^0)$ vs. WIMP-nucleon spin-independent scattering cross-section plane overlain with the limit from LUX-ZEPLIN [181]. (c) and (d) show the $m(\tilde{\chi}_1^0)$ vs. WIMP-proton spin-dependent scattering cross-section plane overlain with the limit from PICO-60 [182]. The pMSSM model cross-sections have been scaled by $(\Omega h^2)/0.12$ for a fair comparison with the limits. 119
- 4.45. Models from the Bino-DM scan which pass all non-dark-matter external constraints in the $m(\tilde{\chi}_1^0)$ vs. Ωh^2 plane, coloured by the dominant annihilation mechanism. (a) includes all models, (b) models surviving constraints from direct detection experiments, (c) models surviving ATLAS Run 2, and (d) models surviving both ATLAS Run 2 and constraints from direct detection experiments. 120
- 4.46. Distributions of (a) the WIMP-nucleon spin-independent scattering cross-section before and after LUX-ZEPLIN [181] and ATLAS Run 2, and (b) the WIMP-proton spin-dependent scattering cross-section before and after PICO-60 [182] and ATLAS Run 2. Models from the EWKino scan with a higgsino- or wino-like LSP that pass all non-dark-matter constraints are included. The pMSSM model cross-sections have been scaled by $(\Omega h^2)/0.12$ for a fair comparison with cross-section limits from direct detection experiments. 121
- 4.47. Models from the EWKino scan with a wino- or higgsino-like LSP that pass all non-dark-matter external constraints in the $m(\tilde{\chi}_1^0)$ vs. Ωh^2 plane, coloured with RGB value set to the wino/higgsino/bino fraction of the LSP. (a) includes all models, (b) models surviving constraints from direct detection experiments, (c) models surviving ATLAS Run 2, and (d) models surviving both ATLAS Run 2 and constraints from direct detection experiments. The dashed horizontal line indicates the observed dark matter relic density. 122
- 4.48. Mass-spectrum for four benchmark models with a bino-like LSP which pass all constraints and are not excluded despite having a mass-spectrum within published ATLAS simplified model contours. Produced using PySLHA [167]. 126

4.49. Mass-spectrum for two benchmark models with a higgsino-like LSP which pass all constraints and are not excluded despite having a mass-spectrum within published ATLAS simplified model contours. Produced using PySLHA [167].	127
5.1. Feynman diagram for the direct production of a pair of charginos, followed by decays to the lightest neutralino via emission of a W boson.	132
5.2. The original expected (dashed lines) and observed (full lines) exclusion contours for each of the analyses used in this combination.	133
5.3. Signal point workspaces available for each analysis.	135
5.4. The absolute and relative difference in the combination CL_s values when correlating experimental systematics but not theoretical systematics (Corr) and leaving all systematics uncorrelated (UnCorr). Note the absolute differences are rounded to two decimal places and for clarity some points at low mass are not displayed.	138
5.5. The absolute and relative difference in the combination CL_s values when correlating just experimental systematics (Corr) and correlating both experimental and theoretical systematics (TheoryCorr). Note the absolute differences are rounded to two decimal places and for clarity some points at low mass are not displayed.	138
5.6. The original CL_s values (produced with HistFitter) for 2L0J wave 1 and 2, and the difference for the recalculated values (produced with PyHF).	140
5.7. The original CL_s values (produced with HistFitter) for the 1L and FullHad analyses, and the difference for the recalculated values (produced with PyHF).	141
5.8. Pull plot from an expected exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (300, 100)$ mass point, which combines the 2L0J and 1L analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pulls are shown in red, the original analysis pulls in black.	143

- 5.9. Pull plot from an expected exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (600, 200)$ mass point, which combines the 2L0J, 1L and FullHad analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pulls are shown in red, the original analysis pulls in black. 144
- 5.10. Pull plot from an expected exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (900, 300)$ mass point, which combines the 1L and FullHad analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pulls are shown in red, the original analysis pulls in black. 145
- 5.11. Combined expected exclusion limit overlaid with the exclusion contours from each individual analysis, for a combination performed (a) without correlating systematics uncertainties and (b) with experimental and theoretical systematics correlated. 146
- 5.12. Combined expected exclusion limit overlaid with the expected CL_s values for each signal point (note for clarity CL_s values are not displayed for all available points). (a) shows the results without correlating systematics uncertainties and (b) with experimental and theoretical systematics correlated. 146
- 5.13. The upper limits on the signal strength μ_{SIG} set by the combination (with correlated systematics), and the relative change with respect to using the analysis with the best expected limit at each point with no combination. Note that for clarity the results are not displayed for all available points. 147
- 6.1. Root-mean-square of the difference between reconstructed and true $p_{\text{T}}^{\text{miss}}$ in bins of average number of interactions per bunch crossing, for (a) $t\bar{t}$ and (b) $Z \rightarrow \mu\mu$ events. The lower panel shows the ratio with respect to the *Tight* working point. A hatched band indicates the statistical uncertainty for the *Tight* working point (which is negligible in most bins). 150

- 6.2. Mass spectrum of pMSSM #1 from the EWKino scan in Chapter 4. The masses of the electroweakinos are shown and the dashed lines indicate the various decay modes. Produced using PySLHA [167]. 156
- 6.3. Neural network architecture for METNet_R. Each hidden layer contains 500 nodes. Both input and target features are normalised. 160
- 6.4. Neural network architecture for METNet_C. Each hidden layer contains 500 nodes. 162
- 6.5. Root-mean-square of the difference between the predicted value of $p_{x,y}^{\text{miss}}$ and $p_{x,y}^{\text{miss, True}}$ in bins of $p_T^{\text{miss, True}}$. METNet_R and METNet_C are shown along with current p_T^{miss} working points for (a) $t\bar{t}$ events, (b) VBF $W \rightarrow e\nu$, events. The lower panel shows the ratio with respect to the *Tight* working point, and the hatched band indicates the statistical uncertainty for the *Tight* working point. 164
- 6.6. Root-mean-square of the difference between the predicted value of $p_{x,y}^{\text{miss}}$ and $p_{x,y}^{\text{miss, True}}$ in bins of number of primary vertices. METNet_R and METNet_C are shown along with current p_T^{miss} working points for (a) $t\bar{t}$ events, (b) VBF $Z \rightarrow \mu\mu$ events. The lower panel shows the ratio with respect to the *Tight* working point, and the hatched band indicates the statistical uncertainty for the *Tight* working point. 165
- 6.7. p_T^{miss} distribution for METNet_R, METNet_C, $p_T^{\text{miss, True}}$ and four p_T^{miss} working points, for (a) $t\bar{t}$ and (b) $Z \rightarrow \mu\mu$ events. Events are unweighted. . . 166
- 6.8. Response of the predicted p_T^{miss} for METNet_R, METNet_C, and four current p_T^{miss} working points (a) $t\bar{t}$, (b) $WW \rightarrow \ell\nu\ell\nu$ events. 166
- 6.9. Two p_T^{miss} performance plots for several p_T^{miss} working points, for a pMSSM signal sample that produces large real p_T^{miss} : (a) Unweighted p_T^{miss} distribution, (b) The root-mean-square of the difference between the predicted value of $p_{x,y}^{\text{miss}}$ and $p_{x,y}^{\text{miss, True}}$ in bins of number of primary vertices. 167

6.10. Distributions of p_T^{miss} for two pMSSM signal points, and WW and $Z \rightarrow \mu\mu$ backgrounds, for (a) <i>Loose</i> p_T^{miss} , (b) <i>Tenacious</i> p_T^{miss} , (c) METNet _R and (d) METNet _C . Events are weighted to 139 fb^{-1} , corresponding to the 2015-2018 ATLAS proton–proton collision dataset luminosity. The lower plot shows signal significance calculated by applying a lower-bound cut at each x -axis bin value.	168
6.11. The average classifier scores and the fraction of events belonging to each class for $t\bar{t}$ and $Z \rightarrow \mu\mu$ simulated test events with 2018 conditions, in bins of (a) average interactions per bunch crossing, (b) number of primary vertices, and (c) sum of the p_T of all objects contributing to the <i>Tight</i> p_T^{miss} working point.	173
6.12. Confusion matrix for the classifier’s prediction of the ‘best’ working point, for (a) $t\bar{t}$ and (b) $Z \rightarrow \mu\mu$ simulated test events with 2018 conditions.	174
A.1. Pull plot from an observed exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (300, 100)$ mass point, which combines the 2L0J and 1L analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pull is shown in red, the original analysis pull in black.	182
A.2. Pull plot from an observed exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (600, 200)$ mass point, which combines the 2L0J, 1L and FullHad analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pull is shown in red, the original analysis pull in black.	183
A.3. Pull plot from an observed exclusion fit for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (900, 300)$ mass point, which combines the 1L and FullHad analyses, with experimental and theoretical systematic uncertainties correlated as outlined in Section 5.3.3. Only systematics where the pull for the combined fit is different to two decimal places compared to one of the original analyses are shown. The combination pull is shown in red, the original analysis pull in black.	184

A.4. Correlations between systematic uncertainties for the $m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0) = (600, 200)$ mass point, for which the 2L0J wave 1, 1L and FullHad analyses are combined.	185
--	-----

List of tables

2.1. The Standard Model particles and their properties. The anti-particles of the quarks and leptons (which have identical mass but opposite charge) are not shown. Mass values are taken from Reference [13].	8
2.2. Field content of the MSSM. Note that only the first of the three quark/squark and lepton/slepton generations are shown for clarity.	14
2.3. The particles beyond-the-Standard Model (except h) introduced by the MSSM. The pre-electroweak symmetry breaking gauge eigenstates and post-electroweak symmetry breaking mass eigenstates are shown. Mixing within the first two generations is assumed to be negligible.	15
2.4. Summary of the 19 parameters of the pMSSM.	21
3.1. Jet selections for ATLAS p_T^{miss} reconstruction working points.	39
4.1. ATLAS searches considered in this analysis	58
4.2. Additional ATLAS results used to exclude pMSSM models.	58
4.3. Filters applied to models before event generation. Models which fail these filters are included in the final results but considered to not be excluded by the direct searches.	61
4.4. Particle-level $\text{CL}_s^{\text{Exp.}}$ criteria used to categorise models. Only models in the ‘ambiguous’ category are run through the detector-level production chain.	64
4.5. Approximations and simplifications made when evaluating each analysis .	68

4.6. Number of models that pass the detector-level selection for each analysis (from Table 4.4), and the number that are actually processed at detector-level. Models with a wino-like LSP were not considered for detector-level production so are omitted from these counts.	69
4.7. Parameter ranges which are randomly sampled for the ‘EWKino’ scan. .	70
4.8. Set-ups for the two pMSSM scans performed and the number of models passing each step. ‘Low statistics truth-production’ here is defined as the number of events generated is less than $2 \times \sigma_{\text{EW}} \times \mathcal{L}$	75
4.9. Simplified model assumptions corresponding to the models for the 2L0J analysis in Figure 4.12.	83
4.10. Simplified model assumptions corresponding to the models for the Compressed analysis in Figure 4.15.	85
4.11. Simplified model assumptions corresponding to the models for the 3L analysis in Figure 4.22.	92
4.12. Additional non-ATLAS constraints from previous collider searches, Higgs measurements, flavour physics observables and electroweak precision measurements.	103
4.13. Percentage of models excluded by each result, for each LSP type and scan.	124
6.1. Simulated event samples with the corresponding matrix element (ME) and parton shower (PS) generators, cross-section order in α_s used to normalise the event yield, underlying-event tune and the generator PDF sets used. .	152
6.2. Summary of object definitions. Baseline objects are used in the $p_{\text{T}}^{\text{miss}}$ calculation while signal objects are used in the event selections.	153
6.3. Jet selections for $p_{\text{T}}^{\text{miss}}$ reconstruction working points used.	154
6.4. Summary of event selections applied when testing METNet performance.	155
6.5. Input features available to METNet. METNet _R uses all these features while METNet _C does not use the x or y components of the $p_{x,y}^{\text{miss}}$ input features or lepton terms.	158
6.6. Neural network hyperparameters.	159

6.7. Summary of the different architectures used for METNet _R and METNet _C .	163
6.8. Maximum signal significance for each working point in Figure 6.10 (plus <i>Tight</i> and <i>TightClus</i>).	169
6.9. Remaining events for different levels of $Z \rightarrow \mu\mu$ rejection using cuts on <i>Tight</i> p_T^{miss} (Tight), regressed METNet _R (Reg.) and the classifier METNet _C (Class.). Values are shown for three signals: $WW \rightarrow \ell\nu\ell\nu$, pMSSM #2 and pMSSM #3.	169