

Searching for Long-Lived Supersymmetric Particles Using Displaced Vertices and
Missing Transverse Energy with the ATLAS Detector

by

Evan Michael Carlson
B.Sc., Pacific University, 2017
M.Sc., University of Victoria, 2019

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Physics and Astronomy

© Evan Michael Carlson, 2024
University of Victoria

All rights reserved. This dissertation may not be reproduced in whole or in part, by
photocopy or other means, without the permission of the author.

Searching for Long-Lived Supersymmetric Particles Using Displaced Vertices and
Missing Transverse Energy with the ATLAS Detector

by

Evan Michael Carlson
B.Sc., Pacific University, 2017
M.Sc., University of Victoria, 2019

Supervisory Committee

Dr. Isabel Trigger, Co-Supervisor
(Department of Physics and Astronomy & TRIUMF)

Dr. Robert Kowalewski, Co-Supervisor
(Department of Physics and Astronomy)

Dr. Scott McIndoe, Outside Member
(Department of Chemistry)

ABSTRACT

The Standard Model of particle physics has been extremely successful in its predictive power and has withstood a wide array of precision tests designed to expose any flaws in its description of fundamental particles. However, the Standard Model is unable to explain several phenomena observed in the universe, such as the nature of the dark matter which makes up more than 80% of the gravitationally interacting matter in the universe. Theories that extend the Standard Model with new fundamental particles have been postulated to address the questions left unanswered by the Standard Model. Many supersymmetric theories provide viable dark matter candidates. In order to more precisely test the Standard Model and its possible extensions, the ATLAS experiment at the Large Hadron Collider has been constructed to measure high energy proton-proton collisions. Long-lived particles (LLPs) are commonly predicted by extensions to the Standard Model. The decay of a LLP to charged particles within the ATLAS Inner Detector would produce tracks that are displaced from the interaction point, which could be reconstructed as a displaced vertex. This dissertation presents a search for displaced vertices with high invariant mass and high track multiplicity in events with significant missing transverse energy in the 2016-2018 data set collected by the ATLAS experiment. The observed number of events is consistent with the number expected from background processes. The results are interpreted in the context of a split-supersymmetry model with long-lived gluinos decaying to neutralinos and Standard Model quarks, and exclusion limits are set at 95% confidence level.

Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	iv
List of Tables	viii
List of Figures	xi
Acknowledgements	xxiii
Dedication	xxvi
Preface	xxvii
1 Introduction	1
2 Theory and Motivation	4
2.1 The Standard Model of Particle Physics	4
2.2 Mathematical Structure of the Standard Model	7
2.2.1 The Electroweak Sector	8
2.2.2 Quantum Chromodynamics	9
2.2.3 Electroweak Symmetry Breaking	10
2.2.4 Long-Lived Particles	12
2.3 Why Search for New Physics?	13
2.4 Supersymmetry	16
2.4.1 The Minimal Supersymmetric Standard Model	17
2.5 Long-Lived Particles in Physics Beyond the Standard Model	19
3 The Large Hadron Collider and ATLAS	22

3.1	CERN and the Large Hadron Collider	22
3.2	Collider Phenomenology	24
3.2.1	Proton-Proton Interactions	25
3.2.2	Anatomy of an LHC Event	27
3.2.3	Luminosity and Cross Sections	29
3.3	The ATLAS Detector	30
3.3.1	Detector Overview and Coordinate System	31
3.4	Magnets	34
3.5	Inner Detector	35
3.5.1	Pixel Detectors	35
3.5.2	Semiconductor Tracker	37
3.5.3	Transition Radiation Tracker	38
3.6	Calorimetry	39
3.6.1	Electromagnetic Calorimeter	40
3.6.2	Hadronic Calorimeter	42
3.7	Muon Spectrometer	42
3.7.1	Precision Tracking Chambers	44
3.7.2	Trigger Chambers	45
3.8	Trigger and Data Acquisition	46
4	Event Reconstruction and Object Definition	49
4.1	Tracking	50
4.1.1	Standard Tracking	50
4.1.2	Large-Radius Tracking	52
4.2	Vertexing	55
4.2.1	Secondary Vertexing	56
4.3	Object Definitions	62
4.3.1	Jets	62
4.3.2	Electrons and Photons	65
4.3.3	Muons	66
4.3.4	Overlap Removal	66
4.3.5	Missing Transverse Energy	67
5	Displaced Vertices + E_T^{miss}: Event and Object Selections	70
5.1	Motivation and Recent DV Searches	71

5.2	Analysis Strategy	73
5.3	Data and Simulated Samples	74
5.3.1	Trigger and DRAW_RPVLL Filter	74
5.3.2	Simulated Samples	75
5.4	Event Selections and Region Definitions	77
5.4.1	Non-Collision Background and Tile Module Vetoes	78
5.4.2	Object Selections	83
5.4.3	Region Definitions	85
5.5	Displaced Vertex Selections	92
5.5.1	Material Veto	94
5.5.2	Track Cleaning	95
5.6	Signal Yields and Efficiencies	100
6	Displaced Vertices + E_T^{miss}: Background Estimation and Uncertainties	107
6.1	Inclusive Background Estimation	108
6.1.1	Jet Matching Method	108
6.1.2	Track Density Method	111
6.1.3	Hybrid Method	116
6.1.4	Uncertainties	121
6.1.5	Validation	126
6.2	Alternative Background Estimation	129
6.2.1	Merged Vertices	132
6.2.2	Accidental Crossings	135
6.2.3	Hadronic Interactions	141
6.2.4	Combined Estimate	142
6.3	Systematic Uncertainties	144
6.3.1	Tracking and Vertexing Uncertainties	144
6.3.2	Jet and E_T^{miss} Uncertainties	146
6.3.3	E_T^{miss} Trigger Uncertainties	147
6.3.4	Theoretical Uncertainties	150
6.3.5	Additional Uncertainties	151
7	Displaced Vertices + E_T^{miss}: Results and Interpretation	153
7.1	Statistical Analysis	155

7.2 Interpretation	159
8 Conclusions	165
Bibliography	169

List of Tables

Table 2.1	The field content of the MSSM [25]. The top three entries are chiral supermultiplets and the bottom three entries are gauge supermultiplets.	18
Table 3.1	Runs of the LHC	23
Table 3.2	Materials, pseudorapidity coverage, and depth of each component of the ATLAS calorimetry system.	43
Table 4.1	Definitions of variables used to parameterize tracks. The reference point used to determine the track’s perigee during track reconstruction is the origin of the detector coordinate system.	50
Table 4.2	Differences between track reconstruction criteria for standard and large-radius tracking [54].	53
Table 4.3	Requirements for standard tracks to be considered for primary vertex construction [55]. The uncertainties on the transverse (d_0) and longitudinal (z_0) impact parameters are denoted $\sigma(d_0)$ and $\sigma(z_0)$	55
Table 4.4	Requirements for standard tracks to be considered in the track jet reconstruction algorithm.	64
Table 4.5	Baseline requirements for EMTopo jets, electrons, photons and muons. Objects passing their respective requirements will be used in the overlap removal procedure. The criteria for the identification working points are discussed in more detail in Section 5.4.2.	68

Table 5.1	Table of lowest unrescaled E_T^{miss} triggers by data-taking period for years 2016, 2017 and 2018. The integrated luminosity \mathcal{L} in pb^{-1} is shown for each period. In the trigger names, the last component denotes the Level 1 trigger used in the trigger chain. For example, a trigger name ending in L1XE55 requires that the event pass the Level 1 trigger requiring 55 GeV of E_T^{miss} . The first component of the trigger name denotes the HLT trigger used in the trigger chain. These components are of the form HLT_xe{X}_{-}{Y}, where {X} is the minimum E_T^{miss} required by algorithm {Y}. Triggers containing an additional _xe{Z} term are required to additionally have at least {Z} GeV of E_T^{miss} as determined by the cell algorithm described in Section 4.3.5. . . .	76
Table 5.2	A summary of the DV mass and track multiplicity cuts for the coarse-binned control, validation and signal regions used in the 1 DV SR inclusive background estimate. These are the primary regions used in estimating and validating the inclusive background estimate. The fine-binned regions in Table 5.3 are only used for diagnostic studies due to limited statistics. . . .	86
Table 5.3	A summary of the DV mass and track multiplicity cuts for the fine-binned control and validation regions used in the 1 DV SR inclusive background estimate. These fine-binned regions are used primarily for diagnostic checks due to low statistics. The coarse-binned regions are the primary regions used for validation. . . .	87
Table 5.4	Summary of the baseline DV selections for the 1 DV SR. . . .	93
Table 5.5	Summary of the signal DV selections for the 1 DV SR. . . .	93
Table 5.6	Selections applied to tracks in DVs reconstructed with VSI. . .	96
Table 6.1	A summary of the uncertainties on the inclusive background estimate for the 1 DV SR.	125
Table 6.2	The estimated values from the extended estimate and the hyper-extended estimate alongside the observed number of DVs in several VRs and the SR (blinded). The selections on m_{DV} and $N_{\text{Tracks}}^{\text{DV}}$ for each region are defined in Table 5.2. The central values of the two estimate methods agree within the statistical uncertainty.	126

Table 6.3	Estimated number of background DVs from the track density and hybrid methods compared to the observed number of DVs in events passing the MTR. Quoted uncertainties for the track density estimate are the full uncertainties discussed in Section 6.1.4. For the hybrid estimate, the uncertainties shown are the non-linearity and pileup uncertainties described in Section 6.1.4 in addition to a statistical uncertainty obtained by assigning an uncertainty of \sqrt{N} to each bin of the histograms used to calculate the estimate.	129
Table 6.4	Definition of radial regions used in the accidental crossings background estimation.	135
Table 6.5	K_S^0 crossing factors for each radial region. A single crossing factor is calculated for the two inner regions due to the lack of statistics in those regions.	140
Table 6.6	Estimated number of background DVs from the track density estimate and the combined background estimate in the validation regions and signal region for the 1 DV SR.	142
Table 6.7	Predicted background in the 1 DV SR from each of the background estimation methods.	144
Table 7.1	Comparison of the estimated number of background events in the 1 DV SR from the various methods described in Chapter 6 to the observed number of events in the 1 DV SR in data. . .	153
Table 7.2	Event-level variables of the 1 DV SR event.	154
Table 7.3	Properties of the EMTopo jets in the 1 DV SR event. The jet η and ϕ are defined with respect to the primary vertex. . . .	154
Table 7.4	Properties of the track jets in the 1 DV SR event. The jet η and ϕ are defined with respect to the primary vertex.	155
Table 7.5	Properties of the DV in the 1 DV SR event.	155
Table 7.6	DV track properties for the DV found in the 1 DV SR.	156

List of Figures

Figure 2.1	The elementary particles that make up the Standard Model of particle physics [3].	6
Figure 2.2	The shape of the Higgs potential (Equation 2.3) in two dimensions. For $\mu^2 > 0$ (a), the potential is symmetric about the minimum and spontaneous symmetry breaking cannot occur, but if $\mu^2 < 0$ (b) the potential is no longer symmetric around any minimum and spontaneous symmetry breaking can occur.	11
Figure 2.3	Mass vs. mean proper lifetime for several fundamental and composite particles in the SM [13]. Particles in the shaded region on the left have lifetimes leading to prompt decays while particles in the shaded region on the right have lifetimes long enough that the majority of their decays will occur outside the ATLAS detector. Particles in the middle region have a significant probability to decay within the ATLAS detector. .	14
Figure 2.4	An example of a top quark loop diagram that provides corrections to the mass of the Higgs boson.	15
Figure 2.5	A Feynman diagram of gluino production and decay in a split-supersymmetry scenario. The \tilde{g} is long-lived and would hadronize into an R -hadron before decaying.	20
Figure 2.6	A Feynman diagram showing associated production of a wino NLSP and a chargino. The wino is long-lived and decays to a bino LSP and a pair of b-quarks via a virtual Higgs boson. .	21
Figure 3.1	A diagram showing the CERN accelerator complex as of 2016 (Run 2) [33]. Among other changes, Linac 2 was replaced with Linac 4 in 2020 in anticipation of Run 3.	25
Figure 3.2	PDF for the proton at energy scale $\mu^2 = 10^4 \text{ GeV}^2$ [34].	26

Figure 3.3	A schematic of a $t\bar{t}h$ event from a pp hard-scatter interaction [35]. The incoming partons are represented by the dark green lines near the center, while the hard interaction is represented by the large red circle in the center. The red lines represent the outgoing particles from the hard-scatter and their resulting radiation and decay products. Blue lines represent particles from the underlying event. A purple oval represents a multiple parton interaction between the partons left over from the hard-scatter. Hadronization is represented by light green ovals, which produce hadrons (dark green circles) that may further decay if unstable. Photons (yellow) can be radiated by charged particles at any stage of the process.	28
Figure 3.4	The mean number of pp interactions per bunch crossing for Run 2 of the LHC [36].	29
Figure 3.5	A cutaway diagram of the ATLAS detector during Runs 1 and 2 of the LHC, with the individual detector subsystems labeled [38].	32
Figure 3.6	A diagram showing how different types of particles leaving the interaction point are detected in the barrel region of ATLAS [39].	33
Figure 3.7	A schematic of the ATLAS magnet system [40]. The barrel toroids are shown in blue, the end-cap toroids in red, and the central solenoid in green.	34
Figure 3.8	A cutaway diagram of the ATLAS Inner Detector [38].	36
Figure 3.9	A quadrant of the ATLAS ID for Run 2 shown as an $R - z$ cross section [43]. The top diagram shows the full ID, while the bottom diagram shows a zoomed in view of the Pixel Detector.	37
Figure 3.10	A cross section diagram of the barrel region of the ID [44]. . .	39
Figure 3.11	A cutaway diagram of the ATLAS calorimeter system [45]. . .	40
Figure 3.12	A cutaway diagram of the ATLAS Muon Spectrometer [38]. . .	44
Figure 3.13	A schematic showing a quadrant of the ATLAS MS [38]. The barrel MDTs are shown in green and the end-cap MDTs are shown in blue. The CSCs are shown in yellow, TGCs are shown in purple, and the RPCs are shown as white boxes.	46

Figure 4.1	A schematic of the parameters used to define tracks in ATLAS [50]. The reference point for track reconstruction is the origin of the detector coordinate system.	51
Figure 4.2	Track reconstruction efficiency as a function of the production radius for displaced charged hadrons produced in the decay of long-lived BSM particles [54]. The red points represent the standard tracking efficiency, the blue points represent the LRT efficiency, and the black points show the sum of the standard and LRT efficiencies.	54
Figure 4.3	An illustration showing the hit pattern requirements imposed on tracks for a vertex decaying close to the pixel B-layer. Because the vertex position is inside the B-layer, the tracks are forbidden from having hits on the IBL and <i>may</i> have hits in the B-layer. In this example, hits are <i>required</i> in Layer 1 of the Pixel Detector [56].	58
Figure 4.4	An example of identifying N -track vertices with the incompatibility graph method. The diagram on the left shows the tracks and vertices in the event, and the diagram on the right shows the incompatibility graph describing the track and vertex composition in the event [57]. Tracks are represented as nodes in the incompatibility graph, and edges represent incompatible pairs of tracks. Compatible vertices are identified by removing nodes until only isolated nodes remain.	59
Figure 4.5	An illustration of how an LLP decaying via b-quarks would be reconstructed by the VSI and FV algorithms [58]. The b-hadron tracks are not reconstructed by the detector due to their short lifetime. Blue circles represent the vertices reconstructed by each algorithm.	62
Figure 5.1	Lower limits on $m_{\tilde{g}}$ as a function of lifetime for (a) fixed $m_{\tilde{\chi}_1^0}$ of 100 GeV and (b) fixed $\Delta m_{\tilde{g},\tilde{\chi}_1^0}$ of 100 GeV [69].	72

Figure 5.2	The E_T^{miss} ϕ distribution for three sample data runs (303304, 303338, 303943). Events are required to pass the DRAW_RPVLL filter selections and have offline E_T^{miss} greater than 150 GeV. The red line shows the distribution with no additional selections applied, and the blue line shows the distribution when applying the SuperLooseBadLLP jet cleaning. The jet cleaning has a negligible effect in suppressing the NCB peaks at 0 and $\pm\pi$.	79
Figure 5.3	The leading jet (a) f_{EM} and (b) f_{max} plotted against the E_T^{miss} ϕ for three sample data runs (303304, 303338, 303943). Events above the horizontal red line are rejected by the NCB veto.	81
Figure 5.4	The leading jet f_{EM} vs. f_{max} for (a) three sample data runs (303304, 303338, 303943) and (b) a gluino sample with $m_{\tilde{g}} = 2200$ GeV, $m_{\tilde{\chi}_1^0} = 100$ GeV, and $\tau = 10$ ns. Events above or to the right of the red lines are rejected by the NCB veto.	82
Figure 5.5	The E_T^{miss} ϕ distribution for three sample data runs (303304, 303338, 303943). The histogram in red shows the distribution with jet cleaning applied, while the histogram in blue shows the distribution when additionally rejecting events which failed the NCB veto. The green histogram shows the distribution after additionally rejecting events which fail the dead tile module veto. The histograms in (a) are normalized to unity to show the effect on the shape of the E_T^{miss} ϕ distribution, while the histograms in (b) are unnormalized to visualize the total reduction in events from applying the vetoes.	84
Figure 5.6	Control regions in the photon-triggered region defined by the DV mass and track multiplicity for (a) the fine-binned regions and (b) the coarse-binned regions.	88
Figure 5.7	Signal and validation regions in the E_T^{miss} -triggered region defined by the DV mass and track multiplicity for (a) the fine-binned regions and (b) the coarse-binned regions.	89
Figure 5.8	Comparison of various event-level variables between the PTR (red) and MTR (blue). All plots are normalized, and the last bin of each histogram contains the overflow.	90

Figure 5.9	Comparison of various jet variables between the PTR (red) and MTR (blue). All plots are normalized, and the last bin of each histogram contains the overflow.	91
Figure 5.10	The number of b-tagged EMTopo jets in the PTR (red) and MTR (blue) in data.	92
Figure 5.11	A map of reconstructed DV positions in the $x - y$ plane for DVs that pass all the baseline requirements in Table 5.4 except the photon veto and max p_T veto [70]. Vertices in (a) have no additional selections applied to them, while vertices in (b) are required to fail the loose material veto.	95
Figure 5.12	Fraction of DV tracks remaining after each track cleaning selection for simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ (black line) and several gluino signal (colored lines) samples. All tracks are required to be associated to a DV which passes the fiducial volume cut, the primary vertex separation cut, and the vertex fit quality cut.	97
Figure 5.13	Track p_T for DV tracks passing all other track cleanings (normalized to unity). The simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples are shown in black, while the colored lines represent simulated gluino signal samples across a range of lifetimes and mass splittings. The p_T of all tracks is shown in (a), while the p_T of attached tracks within and outside the Pixel Detector are shown in (b) and (c), respectively.	98
Figure 5.14	Track $ d_0 /\sigma(d_0)$ for DV tracks passing all other track cleanings (normalized to unity). The simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples are shown in black, while the colored lines represent simulated gluino signal samples across a range of lifetimes and mass splittings. The $ d_0 /\sigma(d_0)$ of all tracks within the beampipe is shown in (a), while the $ d_0 /\sigma(d_0)$ of attached tracks within the Pixel Detector and selected tracks outside the Pixel Detector are shown in (b) and (c), respectively.	99

- Figure 5.15 Track angle with respect to the PV-DV vector ($\Delta\alpha$) for DV tracks passing all other track cleanings (normalized to unity). The simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples are shown in black, while the colored lines represent simulated gluino signal samples across a range of lifetimes and mass splittings. The distribution of $\Delta\alpha$ is shown for (a) all tracks, (b) tracks associated to DVs outside the beampipe, and (c) attached tracks. 101
- Figure 5.16 Efficiency cutflows for simulated (a) gluino R -hadron signal samples with fixed $m_{\tilde{\chi}_1^0} = 100$ GeV and (b) $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples. The signal selections are most efficient for gluinos with a lifetime between 0.1 ns and 3 ns. No simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ events satisfy the signal region selections. The $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples are filtered by the flavor of the jets in the event. 103
- Figure 5.17 Predicted yields (a) and selection efficiencies (b) in the 1 DV SR for gluino samples with a fixed neutralino mass of 100 GeV. The y-axis is $m_{\tilde{g}}$ and the x-axis is the mean proper lifetime τ of the gluino. The expected number of signal events is normalized to an integrated luminosity of 137 fb^{-1} 104
- Figure 5.18 Predicted yields (a) and selection efficiencies (b) in the 1 DV SR for gluino samples with a gluino-neutralino mass difference of 100 GeV. The y-axis is $m_{\tilde{g}}$ and the x-axis is the mean proper lifetime τ of the gluino. The expected number of signal events is normalized to an integrated luminosity of 137 fb^{-1} 105
- Figure 5.19 Predicted yields (a) and selection efficiencies (b) in the 1 DV SR for gluino samples with a gluino-neutralino mass difference of 30 GeV. The y-axis is $m_{\tilde{g}}$ and the x-axis is the mean proper lifetime τ of the gluino. The expected number of signal events is normalized to an integrated luminosity of 137 fb^{-1} 106
- Figure 6.1 ΔR between DVs and the closest track jet in the PTR. DVs are required to pass all baseline selections and have $m_{\text{DV}} > 5$ GeV and $N_{\text{Tracks}}^{\text{DV}} \geq 4$. 23% of DVs in the extended signal region (14 out of 62) have $\Delta R_{\text{DV},\text{Jet}} > 0.4$ 112

Figure 6.2	The number of DVs predicted by the jet matching method in the signal and validation regions. The uncertainties are obtained by propagating the statistical uncertainty on the number of DVs in the PTR used to derive the JDP.	112
Figure 6.3	Histograms used to calculate the EDP in the PTR. Dividing (a) by (b) produces (c). The results shown here are for DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5$ GeV reconstructed outside material. The number of DVs in each bin of (a) are not integers because a pileup weight is applied so that the pileup distribution of the PTR matches the pileup distribution of the MTR. This is discussed in Section 6.1.4.	114
Figure 6.4	Histograms used to calculate the estimated number of DVs in the MTR. Multiplying (a) by (b) produces (c). The results shown here are for DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5$ GeV reconstructed outside material. The total number of estimated DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5$ GeV in the MTR is obtained by integrating (c).	115
Figure 6.5	The number of DVs in each region of the PTR with the full event and DV selections (except the mass and track multiplicity requirements) applied. (a) shows the observed DVs in the fine-binned regions and (b) shows the observed DVs in the coarse-binned regions. The bin contents are not integers because pileup weighting is applied (see Section 6.1.4). The majority of weights are of order 1.	117
Figure 6.6	The estimated number of DVs in the MTR outside material by the (a) hybrid and (b) track density methods.	120
Figure 6.7	The mean and standard deviation of the statistically varied toy experiments for each VR and the 1 DV SR. VR1 ($2 \text{ GeV} < m_{\text{DV}} < 5 \text{ GeV}$, $N_{\text{Tracks}}^{\text{DV}} = 4$) is omitted to allow for better color scaling.	122

Figure 6.8	The fraction of events in the PTR with a SR-like DV ($m_{\text{DV}} > 5 \text{ GeV}$, $N_{\text{Tracks}}^{\text{DV}} \geq 4$) as a function of (a) $p_{\text{T}}^{2, \text{Event}}$, (b) $N_{\text{Track Jets}}^{\text{Event}}$, and (c) $N_{\text{EMTopo Jets}}^{\text{Event}}$. The distribution is fit to a linear function (black line), and the fit parameters are varied within 1σ (red and blue lines) to extract an asymmetric uncertainty on the estimated number of DVs.	123
Figure 6.9	The normalized pileup distribution in the PTR (red) and MTR (blue). A ratio of the MTR to the PTR is shown on the bottom of the plot.	124
Figure 6.10	The number of DVs in the PTR in the (a) extended control region and (b) hyper-extended control region. The values shown are not integers because the DVs are weighted such that the pileup distribution in the PTR matches the distribution in the MTR.	126
Figure 6.11	The estimated (a) and observed (b) number of DVs inside material in the MTR, binned by $N_{\text{Tracks}}^{\text{DV}}$ and m_{DV} . The estimate was performed with the track density method, with the EDP parameterized by $N_{\text{Tracks}}^{\text{Event}}$ and $N_{\text{b-tag}}$	128
Figure 6.12	The observed number of DVs in the MTR, binned by $N_{\text{Tracks}}^{\text{DV}}$ and m_{DV} . The number of DVs in the SR is not shown.	130
Figure 6.13	Summary of the (a) low m_{DV} and (b) low $N_{\text{Tracks}}^{\text{DV}}$ validation regions of the track density estimate. The points in black show the observed number of DVs and the red line shows the value estimated by the track density estimate. The shaded regions show the range of the full uncertainties of the track density estimate.	131
Figure 6.14	Summary of the merged validation regions of the track density estimate. The points in black show the observed number of DVs and the red line shows the value estimated by the track density estimate. The shaded regions show the range of the full uncertainties of the track density estimate.	132

- Figure 6.15 The distance significance squared (S^2) for pairs of DVs in the same event (black) and mixed events (red). DVs are taken from events passing the MTR selections. The mixed-event distribution is normalized such that the integral in the $S^2 > 100$ region is the same for both the mixed-event and same-event distributions. The complete distribution is shown in (a), and (b) shows the distribution in the $S^2 < 100$ region. 134
- Figure 6.16 Normalized mass templates for (a) 4- and (b) 5-track merged vertices. 136
- Figure 6.17 Mass distributions for (a) 2- and (b) 3-choose-2-track DVs for DVs with $R_{xy} < 120$ mm. 139
- Figure 6.18 The estimated background from merged vertices, accidental crossings, and hadronic interactions in signal and validation regions with (a) fine binning and (b) coarse binning. Both plots are binned in m_{DV} and $N_{\text{Tracks}}^{\text{DV}}$ 143
- Figure 6.19 Distribution of the number of K_S^0 candidate DVs as a function of the radial displacement of the DV. The yield in 2016 data is shown in black and the yield in simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown in red. The yield in the simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is normalized such that the yield in the first bin ($R_{xy} < 30$ mm) matches the yield in data. A plot showing the ratio of the yield in 2016 data to the normalized yield in the $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown at the bottom of the figure. 146
- Figure 6.20 Distribution of the number of K_S^0 candidate DVs as a function of the radial displacement of the DV. The yield in 2017 data is shown in black and the yield in simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown in red. The yield in the simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is normalized such that the yield in the first bin ($R_{xy} < 30$ mm) matches the yield in data. A plot showing the ratio of the yield in 2017 data to the normalized yield in the $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown at the bottom of the figure. 147

- Figure 6.21 Distribution of the number of K_S^0 candidate DVs as a function of the radial displacement of the DV. The yield in 2018 data is shown in black and the yield in simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown in red. The yield in the simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is normalized such that the yield in the first bin ($R_{xy} < 30$ mm) matches the yield in data. A plot showing the ratio of the yield in 2018 data to the normalized yield in the $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown at the bottom of the figure. 148
- Figure 6.22 E_T^{miss} trigger efficiencies for several E_T^{miss} triggers used in 2018 data-taking with respect to the offline E_T^{miss} in $Z \rightarrow \mu^+\mu^-$ events in data [68]. Muons are treated as invisible when calculating the offline E_T^{miss} for these events in order to more accurately represent events containing sources of true E_T^{miss} . . . 149
- Figure 6.23 Efficiency of the lowest unprescaled E_T^{miss} trigger active in 2017 with respect to METNoMu for 2017 data (blue) and simulated $Z \rightarrow \mu^+\mu^-$ events (red). The bottom plot shows the ratio of the efficiency in data to the efficiency in simulated events. The agreement in the region used by the analysis ($E_T^{\text{miss}} > 150$ GeV) is within 2%. 150
- Figure 6.24 Efficiency of the lowest unprescaled E_T^{miss} trigger active in 2018 with respect to offline E_T^{miss} for 2018 data (blue) and simulated gluino events (red). The bottom plot shows the ratio of the efficiency in data to the efficiency in simulated events. The maximum deviation between data and simulation in the region used by the analysis ($E_T^{\text{miss}} > 150$ GeV) is less than 3%. . . . 151
- Figure 7.1 Upper limits on the gluino pair production cross section for the split-supersymmetry benchmark model as a function of gluino mass, assuming a fixed mass splitting between the gluino and lightest neutralino of 100 GeV. Limits are set for gluino lifetimes of 0.01 ns (blue), 0.1 ns (orange), 1 ns (red), and 10 ns (brown). The solid black line shows the cross section predicted by the theory. The solid colored lines show the observed upper limit for gluino lifetime, and the dashed colored lines and colored bands show the expected upper limit and its $\pm 1\sigma$ variations. 160

- Figure 7.2 Exclusion limits at 95% confidence level on the mass of the gluino in the split-supersymmetry benchmark model as a function of lifetime for a fixed neutralino mass of 100 GeV. The dashed black line shows the expected limit, and the yellow band shows the area between the $\pm 1\sigma$ variations on the expected limit. The solid red contour shows the observed limit, and the dashed red lines show the observed limits when the gluino pair production cross section is varied up and down by its theoretical uncertainty. 161
- Figure 7.3 Exclusion limits at 95% confidence level on the mass of the gluino in the split-supersymmetry benchmark model as a function of lifetime for a fixed mass splitting between the gluino and neutralino of 100 GeV. The dashed black line shows the expected limit, and the yellow band shows the area between the $\pm 1\sigma$ variations on the expected limit. The solid red contour shows the observed limit, and the dashed red lines show the observed limits when the gluino pair production cross section is varied up and down by its theoretical uncertainty. The area under the curve is excluded. 162
- Figure 7.4 Exclusion limits at 95% confidence level on the mass of the gluino in the split-supersymmetry benchmark model as a function of lifetime for a fixed mass splitting between the gluino and neutralino of 30 GeV. The dashed black line shows the expected limit, and the yellow band shows the area between the $\pm 1\sigma$ variations on the expected limit. The solid red contour shows the observed limit, and the dashed red lines show the observed limits when the gluino pair production cross section is varied up and down by its theoretical uncertainty. The area under the curve is excluded. 163

Figure 7.5 Exclusion limits at 95% confidence level on the mass of the neutralino in the split-supersymmetry benchmark model as a function of gluino mass for a fixed lifetime of 0.1 ns. The dashed black line shows the expected limit, and the yellow band shows the area between the $\pm 1\sigma$ variations on the expected limit. The solid red contour shows the observed limit, and the dashed red lines show the observed limits when the gluino pair production cross section is varied up and down by its theoretical uncertainty. The area under the curve is excluded. 164

ACKNOWLEDGEMENTS

The five years I've spent completing this Ph.D. have been some of the most rewarding and difficult years of my life. I've been extremely lucky to be surrounded by an amazing community throughout my time in graduate school, and this work would never have occurred if it weren't for their support.

To Isabel: you have been the best supervisor I could have possibly asked for. I knew almost nothing about particle physics when you agreed to take me on as your student, and I'm lucky that you were willing to take a chance on me. I've grown so much thanks to your guidance throughout my graduate studies. You've been a constant source of support, encouragement, and wisdom on matters both technical and personal. If there were ever days when you got sick of me constantly knocking on your door, you certainly never let it show. You were always willing to sit and talk through problems with me, even when I was struggling with the same concept for the third time that day. Your rigorous, methodical approach to science and problem-solving has shaped the way I approach problems in all facets of my life. It's been wonderful working with you every day these last five years, and I'm going to miss it.

To Bob: your insightful feedback throughout this whole process has been critical to my success. You were instrumental in getting me on my feet and up to speed when I arrived at UVic, and it would have been a far rockier ride if you hadn't helped me find my way in the early months of my M.Sc.

To Scott: thank you for serving on my committee and providing an outside perspective on topics that particle physicists sometimes take for granted. It's been a pleasure working with you, and I hope you enjoy reading this dissertation.

To the faculty and staff of the UVic Physics and Astronomy department: thank you for your constant support throughout my graduate studies. I'd like to extend special thanks to the faculty and students of the UVic ATLAS group, who have always encouraged me, made me feel welcome and provided any help I asked for.

To the NSW team - Gerardo, Yan, Alam, Estel, Stergios, Artur, Benoit, Denis, Sun, Liang, Camila and many more: thank you for being my family while I was at CERN. You were all so patient with me as I learned more about the NSW project and were wonderful teachers.

To the displaced leptons team: thank you so much for taking me under your wing when I had no idea what to do in an analysis. I know I didn't always make progress as quickly as we wanted to, so I appreciate your patience as I adjusted to working in

an analysis. I'd like to thank Tova, Elodie, and Lesya in particular - you were always so helpful, patient, and encouraging. When I was doubting whether I was actually contributing to the analysis, you were always there to make me feel valued.

To the friends I've made through the TRIUMF ATLAS group - Juan, Leesa, Robin, Dominique, Callum, Damian, Sébastien, Marco, Matt, and many more: I would never have made it through this dissertation without all of you. I've benefitted so much from all the conversations we had in the ATLAS office over the years. Your insight and willingness to talk through problems that were sometimes completely out of your wheelhouse is appreciated more than you could ever know. Thanks for making TRIUMF such a wonderful place to study.

To all the graduate students and postdocs at TRIUMF: thank you for making this such a fun place to work. It's been great having cookies and playing board games with all of you. In particular, thank you to Paul and Daniel - having tea and cocoa with you during the week helped keep me sane and ground me when times were tough.

To Kate: I can't thank you enough for connecting me with the SUSY DV group when I was lost and needed to find an analysis. Since you've been at TRIUMF, you've always found time to talk to me and help me work through problems on the analysis, in spite of your crazy schedule. Our conversations over tea were some of the most enjoyable and productive times I've had at TRIUMF. Your enthusiasm for every inch of progress we fought for in this analysis really helped to keep me going when it seemed like we'd never finish the analysis. Thank you so much for everything.

To Max: thank you for your patient instruction on all things regarding jets and E_T^{miss} . You helped me debug so many key issues throughout this analysis that would have otherwise taken me ages to figure out on my own.

To the ATLAS SUSY/RPVLL group: thank you all for the informative discussions and constructive feedback throughout this analysis. Leading this analysis as a student has been tough, but it was possible thanks to the support of the SUSY community. In particular, I'd like to thank Karri, Evelyn, Javier, and Sara, who have provided many key insights to the team throughout this process.

To the DV + Jets team: when Kate told me I should join the SUSY DV group, she said that the team had some of her favorite people on ATLAS, and now I know why. It was such a pleasure working with you and learning the ins and outs of DV searches. Special thanks go to Christian, Simone, and Emily T., who went above and beyond to mentor me and prepare me to be one of the analysis contacts for DV + E_T^{miss} . I never would have felt like I could do it without your instruction and encouragement.

To the DV + E_T^{miss} team - David R., Risa, Kristin, Jan, Ben, David M., Emily F., and Zubair: it's been an honor to work with you all over the last few years. Thank you so much for your patience and understanding as I grew into my role as co-analysis contact. Doing this analysis has reminded me what I love about physics - working on tough problems with amazing people. You've all been so enthusiastic, insightful, supportive and fun. I couldn't have asked for a better team to grow as a physicist with. I'd like to give special thanks to David Rousso and Ben Rosser, who I had the pleasure to serve alongside as analysis contact. David - it's been a crazy ride, but there's no one else I would have rather driven this analysis with. You've always been there to help, whether it was troubleshooting analysis problems, venting about the roadblocks we'd hit, or sending me memes to help me keep my sanity. Thank you so much. Ben - I'm so thankful that you were willing to step in as a third analysis contact as the analysis moved into the review phase. Your experience in the process and insight into components of the analysis has been invaluable. Thank you for jumping in despite your already tight schedule.

To my church families at Crossroads Thoiry and West Point Grey Baptist Church: you've all been such a blessing to Katie and I during these last few years. There have been some turbulent times over the last few years, but you were a shelter for us through every storm. Thank you for welcoming us into your communities with open arms.

To my in-laws: you've been such a blessing to me since I became a part of this family. Thank you for the constant love and support as Katie and I have bounced back and forth across the globe. We're looking forward to being closer to you soon!

To my family: thank you for giving me all the love and support I could ever ask for. You've always done everything you can to foster my curiosity and love for learning. None of this work would have been possible without you. I love you very much.

And most importantly, to Katie: there aren't any words to express how thankful I am that you came into my life. You've been with me through every up and down since the day we met. When things were at their hardest and I didn't think I could continue, you lifted me up and made sure I knew that everything would be alright. You are a limitless supply of joy and laughter wherever you go. You've been the best partner I could ever imagine. I love you so much, and I can't wait to see what our next adventure brings!

DEDICATION

For Katie, my wonderful wife,
and the family we're building together.

Preface

The ATLAS Collaboration is composed of approximately 6000 physicists, technicians, engineers and students across institutions in 42 countries. The research presented in this dissertation would not have been possible without the involvement of each and every member of the collaboration. Members of the collaboration are responsible for the construction, maintenance and upgrade of the detector, software development, data acquisition, Monte Carlo simulations, and more. It's not possible for one person to contribute to every aspect of the experiment, so the work of each member is critical to the success of the experiment. Furthermore, the research conducted by the collaboration couldn't happen without the successful operation of the Large Hadron Collider.

Considering the collaborative nature of the research presented in this dissertation, it's important to clarify which portions of the work are my direct contributions. The entirety of the text in this document was written by me. Portions of the text appear in or are adapted from the internal ATLAS documentation of the $DV + E_T^{\text{miss}}$ analysis, which I also was a primary author for. All sections of the internal documentation that have been adapted for this dissertation were originally written by me.

Chapters 2, 3 and 4 present background information on the underlying theory of particle physics, the LHC and the ATLAS detector, and the reconstruction algorithms and object definitions needed to understand the analysis that I contributed to. These chapters do not contain any direct contributions from me.

The $DV + E_T^{\text{miss}}$ analysis that I contributed directly to is presented in Chapters 5, 6, and 7. Each analysis in ATLAS has one or more analysis contacts who are responsible for steering the analysis, coordinating the work of the analysis team, and more. From October 2022, I have served as one of the analysis contacts alongside David Rousso and Benjamin Rosser. As part of my role, I have consulted on every analysis decision and study since my appointment. The members of the analysis team are Zubair Bhatti, Oleg Brandt, Kristin Dona, Emily Filmer, Andrew Haas, Paul Jackson, Osamu Jinnouchi, David Miller, Anna Jane Mullin, Jan Offermann, Christian Ohm, Hideyuki Oide, Katherine Pachal, Simone Pagan Griso, Tina Potter, Benjamin Rosser, David Rousso, Isabel Trigger, and Risa Ushioda. My work was done under the supervision of Isabel Trigger with support from the rest of the team and the ATLAS SUSY group. The following list outlines my contributions to the analysis:

- I was responsible for coordinating the production of all simulated background

samples used in this analysis as well as for managing the requests for several extensions to the gluino signal grid targeted by the 1 DV SR.

- I managed the majority of requests for the SUSY15 derivation process needed to process the data and simulated samples, alongside Risa Ushioda.
- I was one of a small group of primary maintainers and developers for the FactoryTools code repository that is used to produce the analysis ntuples. All major ntuple production runs were handled by me. I also ensured that the material map afterburner was run on samples downloaded from the grid before re-uploading the updated versions to the grid so they could be accessed from around the world.
- I validated the object selections and definitions implemented in FactoryTools to ensure they were in line with the recommendations of the various ATLAS combined performance groups.
- I performed detailed studies of the efficiency of the VSI algorithm in order to validate the expected signal yields of the analysis against those of the partial Run 2 DV + E_T^{miss} analysis.
- I performed studies of the code used to implement the DRAW_RPVLL filter in order to ensure that the choice of triggers was consistent across the different data formats utilized by the analysis. This led to the decision to not use 2015 data in the search.
- I conducted studies of the non-collision background contamination in the E_T^{miss} -triggered region and photon-triggered region to develop the loose non-collision background veto used in the 1 DV SR. Risa Ushioda also contributed significantly to this work, including catching the contamination due to dead tile modules and implementing the dead tile module veto.
- I performed the studies comparing the distribution of event-level and jet-level quantities in the PTR and MTR, described in Section 5.4.3.
- A significant amount of the studies used to optimize the event selections in the 1 DV SR were performed by me. These were done in collaboration with David Rousso and Emily Filmer. I also consulted regularly on studies of selections used in the 1 FV and 2 FV SRs alongside David Rousso and Risa Ushioda.

- I performed the studies of the DV track cleaning for VSI DVs, picking up on work previously done by David Rousso and the DV + Jets team.
- I was the primary maintainer for the `plottingscripts` code repository, which is responsible for calculating the expected signal yields and efficiencies for samples in the 1 DV SR. I also validated expected yields in the 1 FV and 2 FV SRs against the yields predicted by David Rousso and Risa Ushioda. The `plottingscripts` repository is also used regularly by other members of the analysis for kinematic studies of signal and background samples.
- I was responsible for the inclusive background estimate for the 1 DV SR. The original jet matching method was developed by Emily Thompson for the DV + Jets analysis. I was responsible for developing and implementing the track density and hybrid methods used in this analysis. The final values predicted by the estimate, the uncertainties on the estimate, and the validation studies are my work. The code I developed was also adapted by Risa Ushioda to perform the inclusive background estimate in the 1 FV SR.
- I regularly consulted with David Rousso on the implementation of the inclusive background estimate for the 2 FV SR.
- I regularly consulted with Kristin Dona to troubleshoot the code used to perform the merged vertex estimate. We worked together often to understand the methods and results of the estimate.
- I worked on troubleshooting the code used to perform the accidental crossings estimate in collaboration with Emily Filmer and Jan Offermann.
- I've had frequent discussions with Jan Offermann regarding the tracking uncertainties in order to troubleshoot the code and understand the results.
- I was responsible for producing the samples used to estimate the jet and E_T^{miss} uncertainties, following the procedures outlined by the relevant ATLAS combined performance groups. I also developed the code in `plottingscripts` used to calculate these uncertainties and pass them to the limit setting framework.
- I worked with Kristin Dona to troubleshoot and develop the code used to estimate the uncertainty from the modeling of the E_T^{miss} trigger turn-on curve in simulation compared to data.

- I requested samples needed to assess the theoretical uncertainties and worked on the code implementation in `plottingscripts` used to calculate those uncertainties and pass them to the limit setting framework.
- I developed the 2D limit plotting framework based on the limit setting implementation developed by Benjamin Rosser. The code I used was adapted from code written by Larry Lee and maintained by Jonathan Long.
- I served as one of the primary authors of the internal documentation for the analysis.

The following figures were produced by other members of the analysis team:

- Figures 6.15, 6.16, 6.23, and 6.24 were produced by Kristin Dona.
- Figures 6.17, 6.19, 6.20, and 6.21 were produced by Jan Offermann.
- Figure 6.18 was produced by Zubair Bhatti, using results from the combined estimate performed by Kristin Dona, Jan Offermann, Emily Filmer, and himself.

All other figures shown in this dissertation that do not have a citation in the caption were made by me. All figures with citations are available under CC-BY licenses or in the public domain, with the exception of Figures 2.3, 3.3, and 4.5. Figure 2.3 originated in [13] and has been reproduced in this dissertation with the permission of Elsevier. Figure 3.3 originated in [35] and has been reproduced here with the permission of Springer Nature BV. Figure 4.5 originated in [58] and has been reproduced here with the permission of the author.

In addition to my involvement in the $DV + E_T^{\text{miss}}$ analysis, I also played a minor role in two previous analyses done by the ATLAS SUSY group, both of which led to publications. Neither of these will be discussed in this dissertation. I was responsible for assessing the uncertainty due to pileup reweighting in the signal samples for a search for displaced leptons. The results of this search are published here:

ATLAS Collaboration, *Search for displaced leptons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Phys. Rev. Lett.* **127** (2020) 051802, arXiv: 2011.07812 [hep-ex]

I also was involved with the $DV + \text{Jets}$ analysis that preceded the $DV + E_T^{\text{miss}}$ analysis that is the focus of my dissertation. I was responsible for performing cross section calculations for the production of fully degenerate higgsinos at high mass.

These cross section calculations were used in the interpretation of the Electroweak RPV scenario in the DV + Jets paper:

ATLAS Collaboration, *Search for long-lived, massive particles in events with displaced vertices and multiple jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **2023.6** (2023) 1, URL: [https://link.springer.com/article/10.1007/JHEP06\(2023\)200](https://link.springer.com/article/10.1007/JHEP06(2023)200)

In addition to analysis work, I was heavily involved in the construction and testing of the ATLAS New Small Wheel upgrade project for the Muon Spectrometer end-cap. I was stationed at CERN from January 2020 through June 2021. During that period, I contributed to the construction of chambers and was primarily responsible for performing x-ray tests of small-strip thin gap chambers (sTGCs) alongside Benoit Lefebvre and Stergios Tsigaridas. The results of these tests were critical to understanding the internal alignment of the strips of the sTGCs. These tests were a natural follow up to the work I performed in my masters thesis, which focused on the results of an sTGC test beam and studies of the internal strip alignment of sTGC chambers.

Chapter 1

Introduction

The universe is a wonderful, mysterious place, and its underlying mechanics have provided endless inspiration to human curiosity throughout history. For millennia, humans have worked to develop a deeper understanding of themselves and the world around them through methodical experiments. This process has led to great strides in our understanding not just of the scientific principles that underlie the universe, but also how humans interact with each other and find meaning in an often turbulent world. Curiosity and exploration have driven this progress throughout history. The work described in this dissertation is a continuation of that tradition.

Pioneering work in quantum mechanics carried out in the 20th century led to the development of the *Standard Model of particle physics* (SM), which describes the interactions of fundamental particles with the strong, weak, and electromagnetic forces. With the discovery of the Higgs boson in 2012 by the ATLAS and CMS experiments, all the fundamental particles predicted by the SM have been discovered. Though the SM has withstood decades of precision tests since its inception, there are fundamental questions about the universe that the SM is unable to answer. More than 80% of the gravitationally interacting matter in the universe does not interact with the electromagnetic force. What is the origin of this so-called ‘dark matter’? What mechanism stabilizes the observed mass of the Higgs boson against large corrections from quantum loops? These are just a few of the questions that motivate high energy physicists to develop a more comprehensive theory of elementary particles. A wealth of theories have been hypothesized that extend the SM with additional fundamental particles in an attempt to answer some of these open questions. One popular class of extensions to the SM is supersymmetry, which extends the particle content of the SM by providing each SM particle with a corresponding superpartner. Many variations of

supersymmetric theories exist, some of which provide viable dark matter candidates and a mechanism to stabilize the Higgs boson mass without fine-tuning.

Dedicated experiments are required to probe the SM and its possible extensions at high precision. Though these experiments can take a variety of forms, high energy particle collisions produced at colliders offer a unique lens with which one can scrutinize the SM. Particle colliders accelerate particles to significant fractions of the speed of light before colliding them inside specialized detectors designed to record the aftermath of the interaction. At high energies, these interactions could produce massive fundamental particles that have yet to be observed. The *Large Hadron Collider (LHC)* is the largest and most energetic collider ever built. Residing 100 m underground on the French-Swiss border near Geneva, the LHC accelerates two beams of protons to a center-of-mass energy of 13.6 TeV¹ in a 27 km ring before colliding them inside state-of-the-art detectors.

The ATLAS detector is one of the detectors at the LHC searching for answers to the questions left open by the SM. The ATLAS Collaboration conducts a wide range of analyses aimed at performing precision measurements of the SM and searching for hints of physics beyond the SM. Extensive searches for supersymmetric particles have been conducted since the start of data-taking in 2009, but no evidence of supersymmetric particles has been uncovered so far. The vast majority of searches for supersymmetry at the LHC have focused on supersymmetric particles with short lifetimes that decay promptly to SM particles. However, there is no a priori reason to believe that supersymmetric particles should have short lifetimes. There are many particles in the SM that have long enough lifetimes that they can travel a measurable distance from their point of production before decaying. These particles are referred to as *long-lived particles (LLPs)*. The decay of a particle is considered *displaced* if the particle travels a discernible distance from the interaction point before decaying. If long-lived supersymmetric particles are being produced at the LHC, then their displaced decays could leave unconventional signatures in ATLAS that could point to their existence. The ATLAS detector and the algorithms it uses to reconstruct particle interactions were primarily designed to analyze promptly decaying particles, so searches for LLP decays require special techniques to analyze the collisions.

This dissertation focuses on a search for long-lived particles that decay inside the tracking volume of the ATLAS detector. These LLPs can decay to SM particles and

¹This dissertation analyzes data collected during Run 2 of the LHC, which had a center-of-mass energy of 13 TeV.

stable particles that escape the detector without interacting. Due to conservation of momentum, the presence of a non-interacting particle produced in an interaction can be inferred from an imbalance of the momentum in the plane transverse to the beamline, referred to as *missing transverse energy* (E_T^{miss}). Tracks from charged SM particles produced in the LLP decay can be traced back to their point of production to form a *displaced vertex* (DV). Because no SM particles produce a DV with high invariant mass and high track multiplicity, searches for DVs have naturally low backgrounds. The analysis detailed in this dissertation searches the data collected by ATLAS between 2016 and 2018 for events containing a DV with high mass and high track multiplicity produced in association with significant E_T^{miss} . This analysis is referred to as the $DV + E_T^{\text{miss}}$ analysis. The results of the analysis are interpreted in the context of a split-supersymmetry model with long-lived gluinos that decay to SM quarks and a stable lightest neutralino, which serves as a dark matter candidate, though the analysis has sensitivity to a variety of theories that predict massive LLPs.

This dissertation is organized as follows. Chapter 2 begins with an introduction to the Standard Model of particle physics and its shortcomings before discussing the motivation for long-lived particle searches. A brief overview of supersymmetry is also presented to provide context for the benchmark split-supersymmetry model used in the interpretation of the results of the analysis. The LHC and the ATLAS detector are described in Chapter 3, and Chapter 4 provides an overview of the reconstruction algorithms used to analyze collisions in ATLAS. Special attention is given to the non-standard reconstruction techniques that are critical to the success of LLP searches in ATLAS. Chapter 5 motivates the $DV + E_T^{\text{miss}}$ analysis and details the selections applied to events and DVs in the data set to optimize the expected sensitivity of the analysis to new physics. The data-driven methods used to estimate the background from detector and algorithmic effects are presented in Chapter 6 alongside a discussion of the systematic uncertainties that are accounted for in the analysis. The results of the analysis are shown in Chapter 7 after a brief introduction to the statistical framework that is used to interpret the results. Chapter 7 also presents an interpretation of the results of the analysis in the context of the split-supersymmetry model described in Chapter 2. Chapter 8 concludes the dissertation with a summary of the results and a discussion of possible future directions for LLP searches at the LHC.

Chapter 2

Theory and Motivation

The Standard Model (SM) of particle physics is the most successful theoretical description of elementary particles and their interactions. This chapter will briefly outline the SM, the particles that compose it, and the mathematical formalism of the theory in Sections 2.1 and 2.2 before discussing a few of the fundamental questions that remain unanswered by the SM in Section 2.3. Section 2.4 outlines the theory of supersymmetry (SUSY), a proposed overarching symmetry that could address some of the questions left open by the SM. The chapter will conclude in Section 2.5 with a discussion of long-lived particles and the unique collider signatures they offer to probe new physics.

2.1 The Standard Model of Particle Physics

Four fundamental forces have been observed in nature - the strong, electromagnetic, weak, and gravitational forces. The SM describes the interactions of elementary particles with all of these forces except gravity. Properties such as the mass, electric charge, weak hypercharge, color charge, and spin are used to characterize these particles and predict how they will interact with each other. All known and hypothesized particles are commonly divided into two groups based on their spin, an intrinsic form of quantized angular momentum carried by elementary particles.

Particles which have a half-integer spin are known as *fermions*, owing to the fact that they obey Fermi-Dirac statistics. Fermions compose all the visible matter of the universe. Each fermion has a corresponding anti-particle with identical spin and mass but opposite charge and *parity*, which describes the behavior of a particle

under spatial inversions. Anti-particles are denoted with a horizontal bar over their symbol, such as \bar{b} . The 12 fermions of the SM are further categorized based on the fundamental forces through which they interact. The six fermions that interact via only the electromagnetic or weak force are known as *leptons* (denoted generically as l). There are three electrically charged leptons - the electron (e), muon (μ), and tau (τ) - each of which has a corresponding neutrino (ν) that carries no electric charge. Each lepton additionally has a *lepton number* L_i , where $i = e, \mu$ or τ . Electrons and electron neutrinos have $L_e = 1$, with the lepton numbers for muons, taus, and their associated neutrinos defined analogously. Anti-leptons have a lepton number of -1 .

The remaining six fermions can additionally interact via the strong force. These particles, known as *quarks* (denoted generically as q), carry *color charge*, which is the analogous form of the electric charge for the strong force. Color charge comes in three types, traditionally denoted as red, blue, and green. Quarks exist in six flavors: up (u), down (d), charm (c), strange (s), top (t), and bottom (b). In units of the proton charge, the up, charm, and top quarks have an electric charge of $+\frac{2}{3}$ and the down, strange, and bottom quarks carry an electric charge of $-\frac{1}{3}$. Quarks with an electric charge of $+\frac{2}{3}$ are collectively called *up-type* quarks and quarks with electric charge of $-\frac{1}{3}$ are called *down-type* quarks. Particles carrying color charge are subject to a phenomenon known as *color confinement*, which requires that these particles only exist in colorless bound states. As a consequence of this, individual quarks are not observed in nature. Instead, quarks are always observed in composite bound states known as *hadrons*. Bound pairs of a quark and antiquark are known as *mesons*. These pairs achieve color-neutrality through the antiquark carrying the anti-color charge of the quark (red and anti-red, for example). A colorless state can also be achieved by a group of three quarks, each with different color charge, known as a *baryon*. Protons and neutrons are both examples of baryons.

Quarks and leptons are each organized into three generations. Each lepton generation consists of a charged lepton and its associated neutrino, and each quark generation consists of a quark carrying electric charge $+\frac{2}{3}$ and a quark carrying electric charge $-\frac{1}{3}$. The lepton generations are ordered with respect to the increasing mass of the charged leptons, while the quark generations are ordered by the quark masses. All charged second and third generation fermions are unstable, so a consequence of this ordering is that charged fermions decay down to their stable first generation counterparts. As a result, all everyday matter is composed of first generation fermions. The division of quarks and leptons into generations is shown in Figure 2.1.




Figure 2.1: The elementary particles that make up the Standard Model of particle physics [3].

One additional property of particles that is important to the mathematical construction of the Standard Model is their chirality. The *helicity* describes the orientation of a particle's spin relative to the direction of the particle's momentum. A particle is said to have right-handed helicity if spin and momentum are aligned and left-handed helicity if the spin and momentum are anti-aligned. However, the helicity of a massive particle is dependent upon the reference frame of the observer, and is thus not Lorentz invariant. *Chirality* is an abstraction of helicity that is Lorentz invariant. The chirality of a particle, which is labelled as right- or left-handed, is important in some fermion interactions [4].

The second group of particles have integer spin and are known as *bosons*. The strong, electromagnetic, and weak forces are mediated by a group of spin-1 particles,

the *gauge bosons*. The mediator of the strong force is the *gluon* (g), which carries color charge, but is massless and electrically neutral. The gluon is responsible for binding quarks into colorless hadrons. The massless *photon* (γ) is the mediator of the electromagnetic force. The weak force is mediated by the W^\pm and Z *bosons*. The Z boson couples to leptons in pairs with a net neutral electric charge (such as $\nu\bar{\nu}$, and e^+e^-) and to same-flavor quark-antiquark pairs (such as $u\bar{u}$). The W^\pm couples to fermions in pairs with a net electric charge of ± 1 (such as $\mu^-\bar{\nu}_\mu$ and $u\bar{d}$). Interactions of the leptons with the Z boson additionally preserve lepton number, meaning interactions such as $Z \rightarrow \mu^-e^+$ are forbidden. Radioactive beta decay of nuclei occurs through the weak force, where the down quark of a neutron converts to an up quark through the emission of a virtual W^- ($d \rightarrow u + W^-$), and the W^- then promptly decays to an electron and electron anti-neutrino ($W^- \rightarrow e + \bar{\nu}_e$). The conversion of the down quark to an up quark changes the neutron into a proton, resulting in a nucleus with an atomic number one greater than the initial nucleus. The last of the bosons is the Higgs boson, a massive spin-0 boson that is responsible for giving mass to most of the particles in the Standard Model. Figure 2.1 summarizes the particle content and organization of the SM [3].

2.2 Mathematical Structure of the Standard Model

The Standard Model is mathematically described by a relativistic quantum field theory. Each particle in the model has a corresponding field, the dynamics of which are described by the SM Lagrangian. Symmetries are key to the mathematical formalism of the SM. The symmetries of the SM arise from the underlying gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$. According to Noether's theorem, each continuous symmetry of a system implies a property of that system that is conserved [5]. For example, symmetry with respect to spatial translations leads to the conservation of linear momentum. The underlying symmetries of the SM give rise to several conserved quantities which are important in the model.

The SM Lagrangian can be divided into several components, each of which describes particular types of particles and interactions:

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{EW}} + \mathcal{L}_{\text{QCD}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}. \quad (2.1)$$

In equation (2.1), \mathcal{L}_{EW} , \mathcal{L}_{QCD} , $\mathcal{L}_{\text{Higgs}}$, and $\mathcal{L}_{\text{Yukawa}}$ are the terms describing electroweak,

strong, Higgs, and Yukawa interactions, respectively. Each of these components will be described in the following sections. This construction results in a renormalizable quantum field theory that is both Lorentz invariant (meaning the laws of physics are identical in all frames of reference, regardless of boost) and gauge symmetric (meaning it is invariant under transformations of the SM gauge group).

2.2.1 The Electroweak Sector

The electromagnetic and weak forces are both the result of a unified *electroweak* theory described by the $SU(2)_L \times U(1)_Y$ gauge group. Electromagnetic interactions between charged particles are mediated by the massless photon and are described by the $U(1)_{EM}$ gauge theory known as quantum electrodynamics (QED). Because photons are massless, the electromagnetic force is a long-range force. All fundamental particles have quantum numbers known as the *electric charge* and *weak hypercharge*, which characterize the interactions of the particle with the electromagnetic and weak forces.

The weak force is mediated by the massive W^\pm and Z bosons and is thus a short-range force. Interactions mediated by the W^\pm are referred to as charged-current interactions and interactions mediated by the Z are referred to as neutral-current interactions. Up- and down-type quarks in the Standard Model are represented as left-handed doublets $(u_L, d_L)^T$ and right-handed singlets u_R and d_R . Each generation of leptons is represented by a left-handed doublet $(l_L, \nu_L)^T$ and a right-handed singlet state l_R . Weak charged-current interactions only couple to left-handed fermions, while weak neutral currents couple to both right- and left-handed fermions. However, the right-handed coupling in weak neutral currents is proportional to the particle's electric charge. If right-handed neutrinos were to exist, they would not interact with either the W or Z bosons.

The theory of electromagnetic and weak interactions can be unified by extending the $U(1)_{EM}$ gauge theory of QED to a $SU(2)_L \times U(1)_Y$ gauge theory, where the subscript L denotes interactions only with left-handed particles and the subscript Y denotes the weak hypercharge [6, 7]. Electroweak interactions are then characterized by four massless fields - the three weak isospin fields W_μ^i ($i = 1, 2, 3$) and the single weak hypercharge field B_μ . The three generators of the $SU(2)_L$ gauge group, denoted T_a , are zero for right-handed fermions. For left-handed fermions, the generators T_a are a set of complex 2×2 matrices with a set of eigenvalues referred to as the weak

isospin. The weak hypercharge, Y , is the generator of the $U(1)_Y$ gauge group. Under this unified electroweak theory, a particle's electric charge under the $U(1)_{\text{EM}}$ gauge group is given by

$$Q_{\text{EM}} = \frac{1}{2}Y + T_3, \quad (2.2)$$

where Y is the weak hypercharge and T_3 is the third component of the weak isospin. Q_{EM} , Y , and T_3 are all conserved in electroweak interactions.

The weak isospin and weak hypercharge fields are not the same as the physical W^\pm , Z and γ particles that mediate the electroweak force, in particular because the W^\pm and Z have mass. The physical states and their mass (or lack thereof) arise from the spontaneous breaking of electroweak symmetry, discussed in Section 2.2.3.

2.2.2 Quantum Chromodynamics

The strong force acts on all particles that carry color charge and is described by the theory of quantum chromodynamics (QCD). Symmetries in QCD arise from the $SU(3)_C$ gauge group, where the subscript C represents the color charge. The eight massless gluons that mediate the strong force correspond to the eight complex 3×3 matrices that are the generators of $SU(3)$ [8]. $SU(3)_C$ is a non-Abelian gauge group, meaning that not all elements of the group commute. Physically, this results in the self-interaction of gluons. As mentioned in Section 2.1, particles carrying color charge are subject to a phenomenon called color confinement, which forbids colored particles from existing outside color-neutral bound states called *hadrons*. Confinement results in interesting phenomena when colored particles are forced apart in high energy interactions. The binding potential between a quark-antiquark pair is found to increase as the particles separate beyond a distance of approximately 0.5 fm. As the particles separate further, it eventually becomes energetically favorable to produce an additional quark-antiquark pair from the vacuum instead of continuing to increase the distance between the bound quark and antiquark. This process, known as *hadronization*, can occur multiple times from the separation of one initial quark-antiquark pair, resulting in a stream of hadrons from the point of the initial interaction. These streams of hadrons are referred to as *jets*. Jets are a key feature of the data at hadron colliders. Jets are discussed further in Section 4.3.1.

2.2.3 Electroweak Symmetry Breaking

The charged fermions and the W^\pm and Z bosons have all been observed to have non-negligible mass. This necessitates the addition of mass terms to the SM Lagrangian. However, explicit mass terms for chiral fermions and vector bosons would violate gauge invariance and thus cannot be included in the SM Lagrangian. The vector bosons instead get their mass through the Brout-Englert-Higgs mechanism (or just “the Higgs mechanism”), which spontaneously breaks the electroweak symmetry of the $SU(2)_L \times U(1)_Y$ gauge group [9, 10].

The spontaneous breaking of electroweak symmetry is achieved through the introduction of a complex scalar doublet field ϕ , which adds to the SM Lagrangian a potential

$$V(\phi) = \mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2, \quad (2.3)$$

where μ and λ are coupling parameters. The shape of this potential is shown in Figure 2.2 for (a) positive and (b) negative values of μ^2 . If μ^2 is positive, the potential has a unique minimum, and it is symmetric about that minimum. In this situation, spontaneous symmetry breaking cannot occur. However, if μ^2 is negative, the potential no longer has a unique minimum, and the potential is not symmetric about any of the possible minima. This scenario results in a ring of minima at radius $\sqrt{\frac{-\mu^2}{2\lambda}} \equiv \frac{v}{\sqrt{2}}$, where v is the vacuum expectation value of the field ϕ .

When the Higgs field assumes a non-zero vacuum expectation value, the electroweak symmetry of $SU(2)_L \times U(1)_Y$ is spontaneously broken, and the weak isospin and weak hypercharge fields can mix to produce the observed electromagnetic and weak interactions. The massless photon and massive Z boson arise from the mixing of the W_μ^3 and B_μ fields, given respectively by

$$A_\mu = \cos(\theta_W) B_\mu + \sin(\theta_W) W_\mu^3 \quad (2.4)$$

and

$$Z_\mu = -\sin(\theta_W) B_\mu + \cos(\theta_W) W_\mu^3, \quad (2.5)$$

where θ_W is the weak mixing angle. The physical states of the W^\pm are given by

$$W^\pm = \frac{1}{\sqrt{2}} (W^1 \mp iW^2). \quad (2.6)$$




Figure 2.2: The shape of the Higgs potential (Equation 2.3) in two dimensions. For $\mu^2 > 0$ (a), the potential is symmetric about the minimum and spontaneous symmetry breaking cannot occur, but if $\mu^2 < 0$ (b) the potential is no longer symmetric around any minimum and spontaneous symmetry breaking can occur.

Electroweak symmetry breaking in the SM also allows for the addition of gauge invariant Yukawa terms to the Lagrangian. The Yukawa terms of the SM Lagrangian describe the interactions of fermions directly with the Higgs field. The coupling terms include both left- and right-handed fermion fields and their coupling strength is proportional to the observed mass of the fermion. These interactions give mass to all charged fermions in the SM.

In addition to imparting mass to the fermions and gauge bosons of the SM, the Higgs mechanism results in a scalar boson as a physical manifestation of the Higgs field, known as the Higgs boson. The ATLAS and CMS experiments at the Large Hadron Collider announced the discovery of a scalar particle with a mass of 125 GeV consistent with the SM Higgs boson in 2012 [11, 12]. Further studies of this particle continue to show that its properties are consistent with the SM Higgs boson. With this discovery, no more particles predicted by the SM were left undiscovered and the model was completed.

2.2.4 Long-Lived Particles

Unless they are prevented from doing so by the conservation laws of the SM, all fundamental particles will spontaneously decay into lighter particles over time. For example, the muon decays through the process $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$. This decay process conserves energy, electric charge, and lepton number, among other symmetries. The electron, on the other hand, cannot decay because it is the lightest charged particle in the SM. Any conceivable decay to a lighter particle would violate conservation of electric charge. Particles that cannot decay to lighter particles are referred to as *stable particles*, and all other particles are considered *unstable*.

The *decay rate* of a particle, Γ , is the probability per unit time that a particle will decay. The value of the decay rate for a particle X to n other particles can be calculated as

$$\Gamma \propto \frac{1}{2m_X} \int |\mathcal{M}|^2 d\Phi_n(p_X; p_1, \dots, p_n), \quad (2.7)$$

where m_X and p_X are the mass and four-momentum of particle X , \mathcal{M} is the matrix element, $d\Phi_n$ is the phase-space factor for the decay, and p_i are the four-momenta of the decay products. The *mean proper lifetime* τ of the particle is the inverse of the decay rate. Considering Equation 2.7 and the relationship of the lifetime and decay rate, a large lifetime can occur if the decay rate is reduced by one of the following mechanisms:

- If the coupling factor g between the particle and its decay products is small, the magnitude of the matrix element \mathcal{M} is reduced because $\mathcal{M} \propto g$.
- If the available phase space for the decay products is small, the decay rate is reduced. This is the case for decays where the mass splitting between parent and daughter particles is small.
- If the decay is mediated by a massive particle, the propagator for the mediating particle suppresses \mathcal{M} .

For a particle with mean proper lifetime τ , the probability that it will survive in the lab frame for a length of time t before decaying is given by

$$P(t) = e^{-t/(\gamma\tau)}, \quad (2.8)$$

where γ is the Lorentz factor of the particle in the lab frame.

A plot showing the masses and lifetimes of several fundamental and composite particles of the SM is shown in Figure 2.3. From an experimental viewpoint, it is helpful to sort particles into one of three categories determined by the spatial resolution of the detector:

- **Detector-prompt** particles are those with a short enough lifetime that the distance between their points of production and decay is typically smaller than the spatial resolution of the detector. Particles such as the W , Z , and Higgs bosons fall into this category. The shaded region on the left of Figure 2.3 contains these particles in the context of the ATLAS detector.
- **Detector-stable** particles are those whose lifetimes are long enough that they typically travel distances much larger than the size of the detector before decaying. Particles such as muons and neutrons are considered detector-stable, and particles such as protons and electrons (which have never been observed to decay) are considered stable. These particles are shown in the shaded region on the right of Figure 2.3.
- **Intermediate lifetime** particles are those with lifetimes such that a significant fraction of those produced in the detector will decay within the volume of the detector. This range of lifetimes is shown in the white central region of Figure 2.3, and includes particles such as the b and K_S^0 mesons and the τ lepton.

Any particle that is not detector-prompt is considered a long-lived particle (LLP). LLPs that decay inside the detector can provide unique signatures that can be studied in relatively low-background environments.

2.3 Why Search for New Physics?

Since its inception, the SM has provided extraordinarily accurate predictions about the nature of fundamental particles and their interactions. Several particles were predicted long before their discovery by experiments, including the W [14, 15], Z [16, 17], and Higgs [11, 12] bosons, as well as the top [18, 19] and charm quark [20, 21]. The predictions of the SM have been validated against numerous high precision tests across a wide range of energies. Despite its success, there are several fundamental

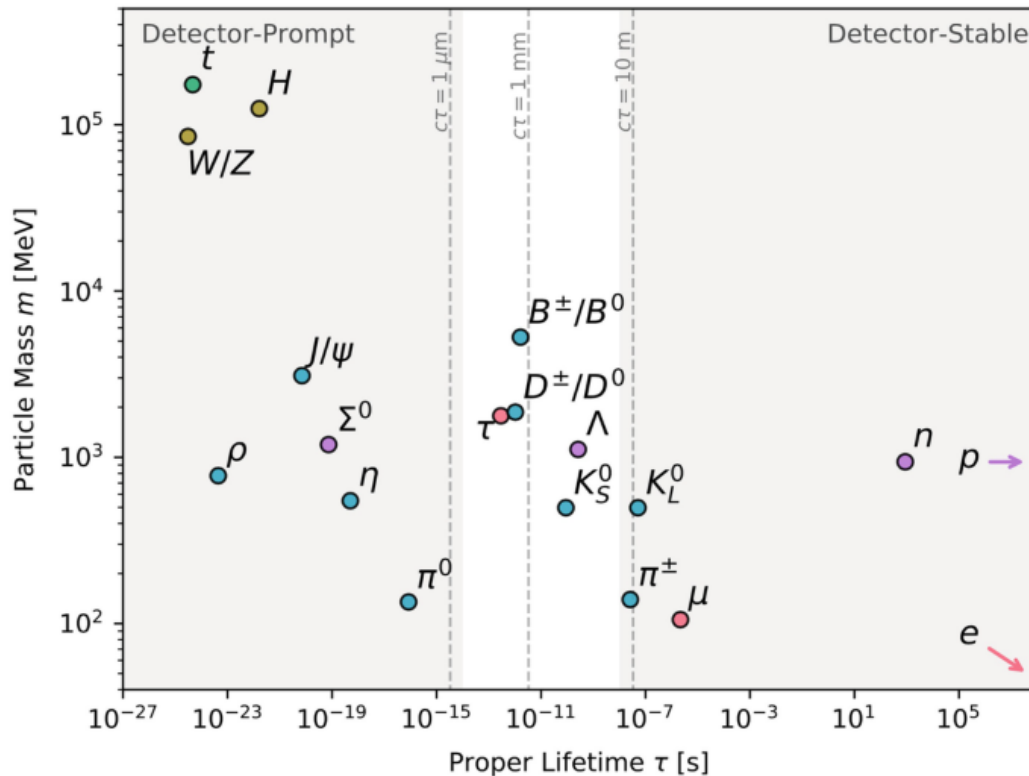


Figure 2.3: Mass vs. mean proper lifetime for several fundamental and composite particles in the SM [13]. Particles in the shaded region on the left have lifetimes leading to prompt decays while particles in the shaded region on the right have lifetimes long enough that the majority of their decays will occur outside the ATLAS detector. Particles in the middle region have a significant probability to decay within the ATLAS detector.

questions from experimental observations that the SM is unable to explain. Two of these questions will be briefly discussed in the remainder of this section.

The physical mass of the Higgs boson is experimentally observed to be 125.11 GeV, which is on the order of the electroweak scale [22]. When calculating the mass of the SM Higgs boson, quantum loop diagrams such as the one in Figure 2.4 must be included in addition to the lowest order diagram. The observed mass of the Higgs can be expressed as

$$m_h^2 = m_{h,0}^2 + \Delta m_h^2, \quad (2.9)$$

where the $m_{h,0}^2$ is the term obtained from the lowest order diagram (known as the *bare* Higgs mass) and Δm_h^2 is the correction term from the loop diagrams. The loop

corrections are proportional to the ultraviolet cutoff scale, Λ , which is the threshold at which the behavior of the model could change significantly due to the influence of potential new physics. The Planck mass, $M_P \approx 10^{19}$ GeV, is the energy at which the gravitational force becomes comparable in strength to the other forces of the SM. If the value of the cutoff scale Λ is set to the Planck mass, the correction term is then proportional to M_P^2 . With the observed mass of the Higgs boson being 125.11 GeV, the value of the bare Higgs mass must be fine-tuned to almost exactly cancel the loop corrections. There is no mechanism in the SM that forbids fine-tuning, but this requirement raises suspicions regarding the “naturalness” of the SM. It is possible that the inclusion of undiscovered particles with couplings to the Higgs could cancel out the loop corrections in a way that removes the need for this fine-tuning. This problem is known as *the hierarchy problem*.




Figure 2.4: An example of a top quark loop diagram that provides corrections to the mass of the Higgs boson.

Another shortcoming of the SM is its inability to explain dark matter. Evidence from astrophysical observations suggests that a significant portion of the gravitationally interacting matter ($\sim 80\%$ by mass) in the universe does not interact with the electromagnetic force. Matter which does not interact with light is invisible, hence the term *dark matter*. The first convincing astrophysical evidence for the existence of dark matter came from the observation of galactic rotation curves, where it is observed that the distribution of the velocities of stars as a function of the radial distance from the galactic center is inconsistent with the observed distribution of the luminous mass of the galaxy [23]. Additional evidence from observations of the cosmic microwave background and gravitational lensing reinforce the need for dark matter in the universe to explain experimental observations.

Currently, the exact nature of dark matter is unknown. The predominant theory is that dark matter is composed of one or more fundamental particles. Any SM dark matter candidate must be electrically neutral and massive, leaving neutral baryonic matter and neutrinos as potential candidates. Astrophysical observations of the elemental abundances of the universe and the large scale structure of the universe have excluded both possibilities. If dark matter were to be a new fundamental particle, it would need to be massive, at most weakly interacting with the known particles of the SM, and stable on cosmological time scales. Despite current constraints, the mass spectrum for viable dark matter candidates spans many orders of magnitude. This has led to a rich and varied experimental program dedicated to determining the true nature of dark matter [24].

2.4 Supersymmetry

Many theories have been postulated in an attempt to address the shortcomings of the SM discussed in the previous section. The vast majority of these theories rely on the introduction of new particles to the SM in order to resolve the deficiencies of the model. These theories are collectively referred to as physics beyond the Standard Model (BSM). Though the search presented in this dissertation is model-independent, it is instructive to discuss an example model of BSM physics that could produce the displaced vertex signature targeted by this analysis.

One prominent class of BSM theories is *supersymmetry* (SUSY). Though there are many variations of SUSY, the core concept of all SUSY models is that an additional symmetry gives every boson a fermionic counterpart and every fermion a bosonic counterpart. All quantum numbers for these *superpartners* remain the same as for their SM partners except for the spin. The superpartner of a particle is denoted with a tilde (\sim) over the symbol of the SM particle. The introduction of this additional symmetry offers an elegant solution to the hierarchy problem. Because there is a relative minus sign between the Higgs mass correction terms from bosons and fermions, the corrections to the Higgs mass from SM particles can be exactly cancelled by the loop corrections from their corresponding superpartner in the case where the superpartners also have identical mass. In addition to addressing the hierarchy problem, many SUSY models provide viable dark matter candidates and some offer mechanisms for the unification of the electroweak and strong forces at high energy scales [25]. Due to its potential to answer many of the largest open questions in particle physics, SUSY is an active topic

of research at the Large Hadron Collider.

2.4.1 The Minimal Supersymmetric Standard Model

The Minimal Supersymmetric Standard Model (MSSM) extends the SM by adding boson-fermion supersymmetry [25]. Each SM field has an associated superpartner field that together form *supermultiplets*. The SM gauge bosons are grouped with their fermionic superpartners into *gauge supermultiplets* [25]. The names of these superpartners are denoted by appending the suffix ‘-ino’ to the corresponding SM gauge boson’s name. For example, the superpartner of the SM gluon is the *gluino*. These superpartners are collectively referred to as *gauginos*.

A *chiral supermultiplet* consists of a SM fermion and two scalar superpartners - one partner for the left-handed fermions and one for the right-handed fermions. The superpartners residing in chiral supermultiplets are conventionally referred to by adding the prefix ‘s’ to the name of their corresponding SM fermion partner. For example, the superpartner of the electron is known as the *selectron*. The superpartners of SM quarks and leptons are collectively referred to as *squarks* and *sleptons*.

If one were to construct the superpartner to the SM Higgs doublet field, it would lead to a violation of gauge invariance in the SUSY sector [25]. The simplest solution to this in the context of the MSSM is to extend the SM Higgs field to instead consist of two complex scalar doublets, H_u and H_d , with a total of eight degrees of freedom. Following electroweak symmetry breaking, three physical degrees of freedom are consumed in providing mass to the W^\pm and Z bosons (as in the SM), with the five remaining degrees of freedom resulting in five physical Higgs bosons - h , H , A , H^+ , and H^- . h and H are neutral scalar bosons (with H more massive than h by convention), A is a neutral pseudoscalar boson, and H^\pm are charged scalar bosons. The Higgs boson discovered at the LHC with a mass of 125 GeV is assumed to correspond to the lightest neutral Higgs boson, h . The fermionic superpartners of this extended Higgs sector, \tilde{H}_u and \tilde{H}_d are referred to as *higgsinos*.

The mixing of the higgsinos with the superpartners of the W_μ^i and B_μ fields (the wino and bino fields, \tilde{W}_μ^i and \tilde{B}_μ) results in two charged mass eigenstates known as charginos and four neutral mass eigenstates known as neutralinos [25]. The two charginos are denoted as $\tilde{\chi}_i^\pm$ and the four neutralinos are denoted as $\tilde{\chi}_i^0$. The mass eigenstate is identified with the subscript i , where $i = 1$ is the lightest mass eigenstate. The field content of the MSSM is summarized in Table 2.1.

Names	Spin 0	Spin 1/2	Spin 1
Squarks, quarks (x3 families)	$(\tilde{u}_L, \tilde{d}_L)$	(u_L, d_L)	-
	\tilde{u}_R^*	u_R^\dagger	-
	\tilde{d}_R^*	d_R^\dagger	-
Sleptons, leptons (x3 families)	$(\tilde{\nu}_L, \tilde{e}_L)$	(ν_L, e_L)	-
	\tilde{e}_R^*	e_R^\dagger	-
Higgs, higgsinos	(H_u^+, H_u^0)	$(\tilde{H}_u^+, \tilde{H}_u^0)$	-
	(H_d^0, H_d^-)	$(\tilde{H}_d^0, \tilde{H}_d^-)$	-
Gluino, gluon	-	\tilde{g}	g
Winos, W bosons	-	$\tilde{W}^\pm \tilde{W}^0$	$W^\pm W^0$
Binos, B boson	-	\tilde{B}^0	B^0

Table 2.1: The field content of the MSSM [25]. The top three entries are chiral supermultiplets and the bottom three entries are gauge supermultiplets.

Unlike in the SM, it is possible to include renormalizable terms in the MSSM that do not conserve baryon number (B) or lepton number (L). Although B and L are not fundamental symmetries of the SM, no processes that violate B or L conservation have been observed experimentally. In order to avoid the introduction of such terms to the MSSM, an additional symmetry known as R -parity can be introduced. The R -parity of a particle is defined as

$$P_R = (-1)^{3(B-L)+2s}, \quad (2.10)$$

where s is the spin of the particle. A consequence of this definition is that all supersymmetric particles (collectively called *sparticles*) have $P_R = -1$ and all SM particles, including the extended Higgs sector, have $P_R = 1$. R -parity is a multiplicative quantum number, meaning that in any given interaction, the product of P_R of the interacting particles must be the same before and after the interaction. SUSY models that require the conservation of R -parity have several important phenomenological characteristics:

- The lightest sparticle (referred to as the *LSP*) must be stable. If the LSP is electrically neutral, it can only interact weakly with SM particles, making it a viable dark matter candidate.
- All sparticles other than the LSP must decay to a state containing an odd

number of LSPs.

- Sparticles can only be produced in even numbers at collider experiments.

These characteristics make R -parity conserving (RPC) models of supersymmetry popular targets for BSM searches at collider experiments. However, no supersymmetric particles have been discovered as of the time of writing. In the MSSM as it has been presented to this point, the superpartners of the SM particles have the same mass as their SM partner. If the MSSM were an accurate description of reality, these superpartners would have been experimentally observed at particle colliders. As such, if the MSSM is to remain consistent with experimental observations, supersymmetry must be a spontaneously broken symmetry. The introduction of explicit *soft supersymmetry breaking* terms to the MSSM can increase the masses of the superpartners significantly [25]. Supersymmetry breaking only occurs in these terms, leaving the high energy behavior of the theory unchanged. The most important consequence of soft supersymmetry breaking is that the cancellation of the loop corrections to the Higgs mass is maintained in the case that the superpartner masses are not too large. This retains the attractive solution to the hierarchy problem.

2.5 Long-Lived Particles in Physics Beyond the Standard Model

Though there are many SUSY scenarios that predict long-lived particles, the majority of searches for supersymmetric particles at the LHC have focused on signatures involving promptly-decaying particles. As expanding regions of the parameter space have been excluded by prompt searches, however, an increasing amount of focus has been dedicated to searching for long-lived particles predicted by various BSM scenarios. Though the search for displaced vertices detailed in this dissertation is signature-driven and model-independent, two SUSY models will be presented in the remainder of this section to motivate the displaced vertex signature. Outside of SUSY, many models of BSM physics, such as Higgs portal and hidden valley models of dark matter, predict long-lived particles capable of producing *displaced vertices* (DVs) [26, 27].

SUSY models that successfully address the hierarchy problem are heavily constrained by the so-far negative results of searches for supersymmetric particles at the LHC. Models of split-supersymmetry choose to relax the requirement on naturalness

in favor of a model that retains other attractive properties - namely a viable dark matter candidate and the unification of the strong and electroweak forces at high energy [28]. In split-supersymmetry, the superpartners of the SM fermions (squarks and sleptons) have masses near or above the supersymmetry breaking scale, while the masses of the gauginos and higgsinos are significantly lighter, on the order of 1 TeV. Though the squarks and sleptons in split-supersymmetry are inaccessible at LHC energies, gauginos and higgsinos could possibly be produced at the LHC. The gluino in split-supersymmetry must decay via a squark to a quark, antiquark and neutralino LSP. However, the squark masses in the model are far greater than the potential gluino mass, which forces the squark in the decay chain to be highly virtual. Because the decay must be mediated by this heavy virtual squark, the gluino obtains a lifetime proportional to $m_{\tilde{q}}^4/m_{\tilde{g}}^5$. Depending on the squark and gluino masses, the gluino lifetime could be arbitrarily long, although lifetimes of less than 100 s are favored due to cosmological constraints [29]. Due to color confinement, gluinos with a lifetime longer than the hadronization timescale (10^{-23} s) will hadronize into *R-hadrons*, which are bound states containing a strongly interacting sparticle and strongly interacting SM particles. Figure 2.5 shows a Feynman diagram of gluino production and decay in a split-supersymmetry scenario. A simplified model of split-supersymmetry where the gluinos become long-lived due to the high mass of their decay mediators serves as a benchmark model for the $DV + E_T^{\text{miss}}$ analysis detailed in this dissertation.




Figure 2.5: A Feynman diagram of gluino production and decay in a split-supersymmetry scenario. The \tilde{g} is long-lived and would hadronize into an *R-hadron* before decaying.

In scenarios where the higgsino mass parameter μ is very large, it is possible for

the next-to-lightest supersymmetric particle (NLSP) to become long-lived [30]. The decay of a neutral wino NLSP ($\tilde{\chi}_2^0$) to a neutral bino LSP ($\tilde{\chi}_1^0$) is suppressed by factors proportional to $1/\mu$ and $1/\mu^2$. If the mass splitting between the wino and bino is smaller than the mass of the Z boson, it is preferential for the decay to proceed through a virtual Higgs boson, further contributing to the suppression of the decay. A Feynman diagram showing such a decay is shown in Figure 2.6. In these scenarios, the lifetime of the neutral wino can result in decays inside the tracking volume of a large detector, such as the ATLAS detector. The production of two displaced b-quarks in such a decay could leave a displaced vertex signature in the detector in addition to missing transverse energy (described in Section 4.3.5).




Figure 2.6: A Feynman diagram showing associated production of a wino NLSP and a chargino. The wino is long-lived and decays to a bino LSP and a pair of b-quarks via a virtual Higgs boson.

Chapter 3

The Large Hadron Collider and ATLAS

To test the Standard Model and other models of BSM physics, physicists have designed complex systems of particle accelerators and detectors to produce and measure fundamental particles with exceptional precision. The search detailed in this dissertation analyzes data collected by the ATLAS detector at CERN's Large Hadron Collider. Section 3.1 will briefly discuss the history of the European Organization for Nuclear Research (CERN) before giving a basic overview of the LHC. A few key principles of hadron collider physics will then be discussed in Section 3.2 to give context to the discussion of the ATLAS detector presented in Sections 3.3 through 3.8. The individual components of the ATLAS detector will be described in Sections 3.4 through 3.7. Section 3.8 will conclude the chapter with an overview of the trigger system used to filter and record data.

3.1 CERN and the Large Hadron Collider

Originally founded in 1954 to research atomic nuclei, CERN is a research organization dedicated to studying subatomic particles and their interactions. Located on the French-Swiss border near Geneva, CERN is now home to a variety of nuclear and particle physics experiments. At the time of writing, CERN consists of 23 Member States throughout Europe, 11 Associate Member States, and two countries with Observer status [31]. Several key discoveries in subatomic physics have been made at CERN, the most recent of which was the discovery of the Higgs boson in 2012 by the

ATLAS and CMS collaborations [11, 12].

CERN is home to the world’s highest energy particle accelerator, the *Large Hadron Collider (LHC)*. First turned on in 2008, the LHC is a 27 km circumference synchrotron accelerator located 100 m underground designed to accelerate two proton beams to a center-of-mass energy (\sqrt{s}) of 14 TeV. These beams are circulated in opposite directions and collided at four points along the accelerator ring at a rate of 40 MHz. Two runs of the LHC have been completed at the time of writing, with a third currently in progress. Table 3.1 lists the different runs of the LHC and their center-of-mass energies. The data used in the analysis described in Chapters 5, 6, and 7 was collected during Run 2 with $\sqrt{s} = 13$ TeV. The following discussions of the LHC and the CERN accelerator complex are representative of the conditions during Run 2.

Run	Dates	\sqrt{s} (TeV)
1	2010-2013	7, 8
2	2015-2018	13
3	2022-2025	13.6
4-6	2029-2041	14

Table 3.1: Runs of the LHC

Protons are accelerated by passing them through radiofrequency cavities (RF) containing oscillating electromagnetic fields [32]. These cavities are tuned to resonate at specific frequencies such that the protons pick up additional energy with each pass through the cavity. Because the direction and magnitude of the field in the cavity vary over time, the protons experience different forces depending on their time of arrival in the cavity. A consequence of this is that a uniform, continuous beam is not achievable when using RF cavities. Instead, protons are grouped into dense *bunches* separated by empty intervals. Each proton bunch at the LHC contains approximately $1.15 \cdot 10^{11}$ protons. A series of superconducting magnets is used to steer the proton beams along their path as well as to focus the beams to tune the collision rate at the interaction points. The LHC steers the proton beams along their circular trajectories with over 1200 dipole magnets, each of which is 15 m long and capable of producing a magnetic field of 8.3 T. The bending of a charged particle’s trajectory in a magnetic field is dependent upon the particle’s momentum, so the strength of the magnetic fields in the dipole magnets increases alongside the proton energy in order to keep the beams on track. Superconducting quadrupole magnets are used to focus the beam at the interaction points, causing the protons to squeeze into a tighter area and interact at a

higher rate when colliding with the oncoming beam.

Accelerating protons to LHC energy scales from rest with a single machine is not feasible, so a chain of injector accelerators is used to gradually increase the proton energy before they enter the LHC. Protons from a bottle of hydrogen gas are stripped of their electrons before being accelerated through the following chain of four accelerators on their way to the LHC:

1. **Linac 2**, a linear accelerator that accelerates the protons from rest up to an energy of 50 MeV¹.
2. The **Proton Synchrotron Booster**, which accelerates protons from Linac 2 to an energy of 1.4 GeV.
3. The **Proton Synchrotron**, which accepts protons from the Proton Synchrotron Booster and accelerates them to an energy of 25 GeV.
4. The **Super Proton Synchrotron (SPS)**, which accelerates protons around a 7 km circumference beamline from 25 GeV to 450 GeV. In addition to serving as an injector to the LHC, the SPS also provides beams to several experiments, including the NA62 and COMPASS experiments.

Figure 3.1 shows the state of the CERN accelerator complex as it was in 2016 [33].

3.2 Collider Phenomenology

The collision of particles at high energy can produce more massive particles that would otherwise be unobservable in nature due to their short lifetime and low natural abundance. This is allowed in the case where the energy of the initial interaction is greater than the rest mass of the new particles produced. The ability to produce massive particles in the heart of specialized detectors makes colliders an excellent environment both for precisely studying the properties of SM particles that cannot be observed directly in nature and for attempting to produce potentially massive undiscovered particles. This section will provide an introduction to some key concepts for studying fundamental particles at hadron colliders.

¹Linac 2 was replaced with Linac 4 in 2020. The description in the text and the diagram in Figure 3.1 are representative of the accelerator complex in 2016 during Run 2.

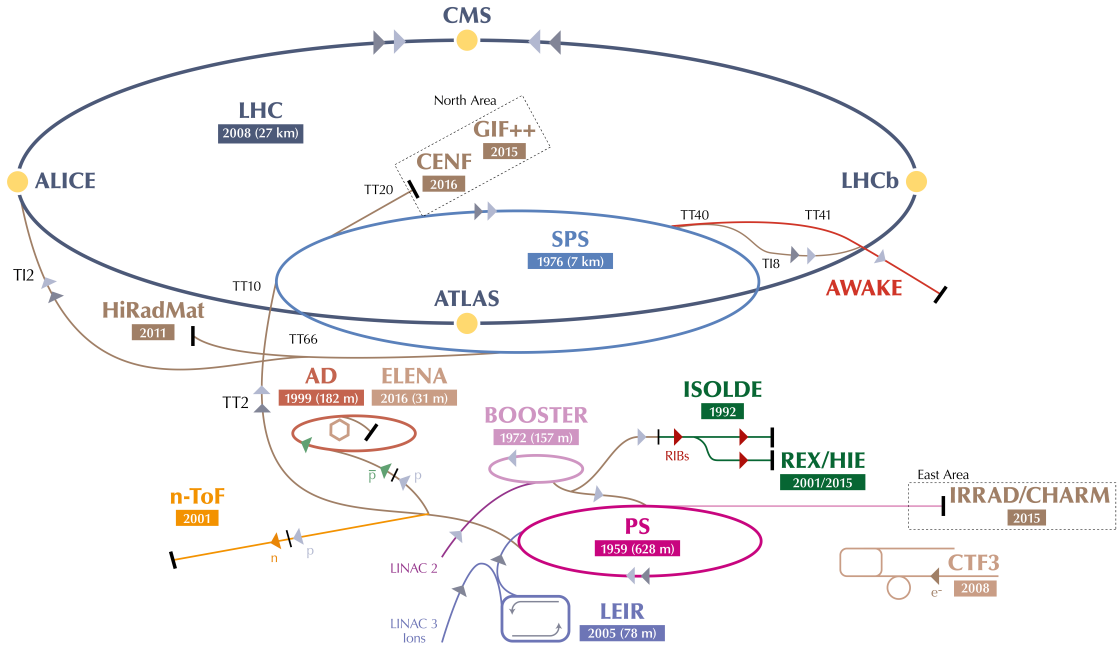


Figure 3.1: A diagram showing the CERN accelerator complex as of 2016 (Run 2) [33]. Among other changes, Linac 2 was replaced with Linac 4 in 2020 in anticipation of Run 3.

3.2.1 Proton-Proton Interactions

As mentioned in Section 2.1, hadrons such as protons are not fundamental particles, but rather composite particles composed of quarks and gluons. The quarks that determine the quantum numbers of a hadron are known as *valence quarks*. In the case of the proton, the valence quarks are two up quarks and a down quark. These valence quarks are bound together through the exchange of gluons. Protons additionally contain an indefinite number of virtual quark-antiquark pairs originating from gluon exchange, which are referred to as *sea quarks*. The valence quarks, sea quarks and gluons in a proton are collectively referred to as *partons*, with each parton carrying a fraction of the total momentum of the proton. In a proton-proton (pp) interaction at LHC energies, it is the individual partons of the proton that interact with the partons of the oncoming proton.

The distribution of the momentum fraction carried by individual partons in a proton is described by a parton distribution function (PDF), $f_a(x, \mu^2)$, where a is the parton type, x is the momentum fraction, and μ^2 is the energy scale at which the proton is being probed. Due to the non-perturbative nature of QCD, the proton PDFs

cannot be calculated analytically and must instead be determined from experimental measurements [34]. A PDF for the proton at energy scale $\mu^2 = 10^4 \text{ GeV}^2$ is shown in Figure 3.2.




Figure 3.2: PDF for the proton at energy scale $\mu^2 = 10^4 \text{ GeV}^2$ [34].

An important consequence of the composite nature of the proton is that although the total momentum of the proton in an LHC collision is well known, the momentum of the individual interacting partons is not known. Even though the initial momentum of the interacting system is unknown, conservation of momentum is still a critical principle when analyzing data in pp interactions. Because momentum is a vector and the initial proton momentum is along the beamline, the longitudinal component of the interacting parton momentum is unknown. However, the proton, and thus the individual partons, carry negligible momentum in the plane transverse to the beamline. This plane is defined by the orthogonal Cartesian coordinates x and y . The vector nature of momentum thus ensures that each individual component of the

initial momentum vector is conserved in the interaction, meaning that the sum of the *transverse momenta* of the outgoing particles from a collision must be zero. A particle's transverse momentum, $p_T = \sqrt{p_x^2 + p_y^2}$, is a key variable in hadron collider physics. Conservation of transverse momentum is used to define a quantity known as *missing transverse energy* (E_T^{miss}) that is essential when searching for new particles that would escape the detector without interacting. E_T^{miss} is a core piece of the analysis described in this dissertation, and it will be discussed further in Section 4.3.5.

3.2.2 Anatomy of an LHC Event

Proton collisions at the LHC produce a wide variety of different processes. The overwhelming majority of these processes are *soft*, meaning that they involve a relatively small momentum exchange. While these soft interactions can be interesting, most of the physics of interest occurring at the LHC comes from *hard-scatter interactions* involving a significant momentum transfer between the interacting partons. The amount of data that can be recorded by the LHC experiments is limited by the various detector technologies and the available computing resources, so identifying and recording the hard-scatter interactions for study is a key challenge for the experiments. A *trigger* is used to identify and record hard-scatter interactions for offline study. The trigger system of the ATLAS experiment will be discussed in more detail in Section 3.8. Each bunch crossing containing at least one pp interaction is referred to as an *event*.

When two partons of colliding protons interact in a hard-scattering interaction, the remaining partons which did not participate in the interaction are no longer bound to the proton. It's possible for these partons to interact with each other, resulting in *multiple parton interactions*, but more commonly the free partons will hadronize and produce relatively soft underlying jets. Additionally, both the incoming and outgoing partons of the interaction can radiate gluons, which will also hadronize and form jets. Jets from gluons radiated before or after the hard-scattering process are referred to as *initial- and final-state radiation (ISR and FSR)*, respectively. The contributions to the event from the proton fragmentation and ISR are collectively referred to as the *underlying event*. Figure 3.3 shows an illustration of a $t\bar{t}h$ event from a pp hard-scatter interaction [35].

In addition to potentially containing a hard-scatter interaction, each bunch crossing at the LHC can contain additional pp interactions. These additional interactions




Figure 3.3: A schematic of a $t\bar{t}h$ event from a pp hard-scatter interaction [35]. The incoming partons are represented by the dark green lines near the center, while the hard interaction is represented by the large red circle in the center. The red lines represent the outgoing particles from the hard-scatter and their resulting radiation and decay products. Blue lines represent particles from the underlying event. A purple oval represents a multiple parton interaction between the partons left over from the hard-scatter. Hadronization is represented by light green ovals, which produce hadrons (dark green circles) that may further decay if unstable. Photons (yellow) can be radiated by charged particles at any stage of the process.

between protons not involved in the hard-scatter interaction are referred to as *pileup*. Pileup interactions generally produce particles with low transverse momenta, but it is possible for multiple interactions with large momentum transfers to occur in a single bunch crossing. The hard-scatter is identified as the process with the largest momentum transfer in these cases. The average number of interactions per bunch crossing varied significantly across Run 2 of the LHC. Events collected in 2015 on average had 13.4 pp interactions per bunch crossing, while events from 2017 had 37.8 pp interactions per bunch crossing on average [36]. Figure 3.4 shows the distribution of the number of pp interactions per bunch crossing for each year of Run 2 as measured by the ATLAS detector. The distributions are weighted to match the integrated

luminosity of the corresponding year. The complexity of individual pp interactions coupled with the number of interactions in each bunch crossing makes the analysis of data at the LHC technically challenging. In order to identify interesting events and extract meaningful measurements from them, elaborate and highly specialized detectors are required to reconstruct the individual components of the interactions with exquisite precision.




Figure 3.4: The mean number of pp interactions per bunch crossing for Run 2 of the LHC [36].

3.2.3 Luminosity and Cross Sections

One way to test the SM and its possible extensions is to compare the number of observed events consistent with a particular process to the number of events the SM predicts for that final state. A necessary component of such a prediction is the probability that the process will take place in a system with a particular initial state. This probability is known as the *cross section* (denoted σ). Cross sections are expressed in units of area, such as cm^2 . Typical cross sections for processes of interest can be extremely small when expressed in cm^2 , so an alternative unit known as the *barn* (b) is more commonly used, where $\text{b} = 10^{-24} \text{cm}^2$. The cross section for

a process can be calculated from the theory and is dependent on the energy of the initial interaction, among other factors. The number of events expected to occur for the process is therefore given by

$$N = \sigma \int L dt, \quad (3.1)$$

where L is the *instantaneous luminosity*. The instantaneous luminosity is a property of the colliding beams that determines the rate of interactions, given by

$$L = \frac{N_1 N_2 f N_b}{2\pi \Sigma_x \Sigma_y}, \quad (3.2)$$

where N_1 and N_2 are the number of protons in each bunch, N_b is the number of bunches, f is the LHC revolution frequency, and $\Sigma_{x,y}$ are the horizontal and vertical convolved beam sizes [37]. Instantaneous luminosity is measured in units of $\text{cm}^{-2}\text{s}^{-1}$ or $\text{b}^{-1}\text{s}^{-1}$. The integral of L over a given time is referred to as the *integrated luminosity*, \mathcal{L} , which can be substituted into Equation 3.1 to simplify the number of expected events to

$$N = \sigma \mathcal{L}. \quad (3.3)$$

The LHC delivered a total integrated luminosity of 157 fb^{-1} to the ATLAS experiment during Run 2, of which 140 fb^{-1} was recorded under good detector conditions and deemed good for physics analysis [37]. The difference between the delivered and usable luminosity arises from inefficiencies in the data acquisition system and the short periods where one or more detector subsystems were not functioning properly.

3.3 The ATLAS Detector

The ATLAS detector is one of two general purpose particle detectors at the LHC² and the largest of the four primary LHC experiments. Located 100 m underground at Point 1 of the LHC, ATLAS is designed to perform precision measurements of the SM and to search for rare or new processes in the LHC data set. Such measurements require quick and precise reconstruction of the momenta, energies, and trajectories of particles produced in the primary interaction. This is accomplished through a suite of specialized detector components, each designed to measure particular types of

²The other general purpose detector at the LHC is the Compact Muon Solenoid (CMS).

particles. ATLAS is composed of four primary detector subsystems:

- A set of superconducting electromagnets used to bend the trajectories of charged particles as they traverse the detector, collectively referred to as the *magnet system*.
- The *Inner Detector (ID)*, whose primary purpose is to reconstruct the trajectories of charged particles produced inside ATLAS. These trajectories, referred to as *tracks*, provide measurements of a particle’s momentum and charge. Track information is critical to the reconstruction of primary and secondary vertices, which will be discussed in detail in Sections 4.1 and 4.2.
- *Calorimeters*, which are used to measure the energies of strongly and/or electromagnetically interacting particles, such as electrons, photons, and hadrons.
- The *Muon Spectrometer (MS)*, which measures the trajectories of muons in a magnetic field using dedicated tracking chambers.

Each of these detector components will be discussed in Sections 3.4 through 3.7.

3.3.1 Detector Overview and Coordinate System

ATLAS is an approximately cylindrical detector with a length of 44 m and a diameter of 25 m [38]. The LHC beamline runs along the length of the detector through the centers of the circular ends. Layers of sub-detectors are arranged around the *interaction point (IP)* to provide hermetic detector coverage for the products of the interaction. The ATLAS geometry can be subdivided into two regions - the barrel and the end-caps. The *barrel* region consists of cylindrical detector layers arranged concentrically around the beampipe, while the *end-caps* are composed of disk-like detector structures oriented in planes transverse to the beamline. Both regions contain detectors from each of the detector subsystems so that ATLAS has complete coverage for all stable particles found in the SM, except for neutrinos, which only interact weakly and thus escape the detector without interacting. A schematic of the ATLAS detector with the individual detector subsystems labeled is shown in Figure 3.5 [38]. Figure 3.6 depicts a cross section of the ATLAS barrel region and how different particle types interact with the individual detector subsystems as they leave the IP [39].

The nominal interaction point serves as the origin of the ATLAS coordinate system. The geometry of the detector is most easily described in a cylindrical coordinate




Figure 3.5: A cutaway diagram of the ATLAS detector during Runs 1 and 2 of the LHC, with the individual detector subsystems labeled [38].

system, with the z -axis defined by the beamline. The xy -plane lies transverse to the beamline, with the positive x -axis pointing to the center of the LHC ring and the positive y -axis pointing towards the sky. The distance between a given point and the IP in the transverse plane is given by $R = \sqrt{x^2 + y^2}$. The azimuthal angle ϕ in the transverse plane is measured from the x -axis such that $x = R \cdot \cos(\phi)$ and $y = R \cdot \sin(\phi)$.

Analysis of ATLAS data is typically performed in a coordinate system defined by two angles - the azimuthal angle, ϕ , and the polar angle, θ . The angle ϕ is defined identically to the cylindrical coordinate system, while θ measures the angle relative the beamline. Because the initial momentum of an interaction along the z -axis is unknown, the outgoing system of particles can have a significant boost along the z -axis relative to the lab frame. The polar angle between two particles produced in a collision is not Lorentz invariant due to relativistic dilation effects, so θ is not typically used to describe particle trajectories. A new quantity, known as *pseudorapidity* (η), is defined as




Figure 3.6: A diagram showing how different types of particles leaving the interaction point are detected in the barrel region of ATLAS [39].

$$\eta \equiv -\ln[\tan(\theta/2)]. \quad (3.4)$$

A convenient feature of pseudorapidity is that the difference in η between two particles is Lorentz invariant to boosts along the z -axis in the limit that the particles are highly relativistic, where it approaches the particle's *rapidity*, $y = \frac{1}{2} \ln((E + p_z)/(E - p_z))$, which is precisely invariant to boosts along the z -axis. It follows from this definition that particles travelling in the transverse plane have $\eta = 0$, while particles travelling along the beamline have $\eta = \pm\infty$. The angular distance between two particles in ATLAS is commonly expressed through the quantity $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ ³.

³Note that ΔR is different from the transverse displacement from the beamline, $R = \sqrt{x^2 + y^2}$.

3.4 Magnets

Charged particles traversing a magnetic field $\vec{\mathbf{B}}$ will experience a force

$$\vec{\mathbf{F}} = q \cdot \vec{\mathbf{v}} \times \vec{\mathbf{B}}, \quad (3.5)$$

where q is the electric charge of the particle and $\vec{\mathbf{v}}$ is the velocity of the particle. The result of this force is that the trajectories of charged particles will bend as they travel through a magnetic field. Because the force a particle experiences is dependent upon the velocity of the particle, the particle's momentum can be determined by measuring the amount of curvature in the particle's trajectory. To take advantage of this fact, the components of ATLAS that measure trajectories are immersed in strong magnetic fields. The magnetic fields in ATLAS are provided by three subsystems - the central solenoid, the barrel toroids, and the end-cap toroids. A diagram of the magnet system is shown in Figure 3.7 [40]. All three magnet subsystems produce their fields with thousands of coils of superconducting Nb-Ti wire. The coils must be continuously cooled to a temperature of 4.5 K to maintain their superconducting properties [38].




Figure 3.7: A schematic of the ATLAS magnet system [40]. The barrel toroids are shown in blue, the end-cap toroids in red, and the central solenoid in green.

The central solenoid contains the entirety of the Inner Detector and provides a 2 T magnetic field parallel to the beamline. Charged particles traversing the Inner Detector have their trajectories bent in the transverse plane. This bending is used to measure the transverse momentum, p_T , of the particle. Located outside the calorimeters, the barrel and end-cap toroids each consist of eight toroidal coils of superconducting wire, spaced evenly in the 360° of ϕ around the beamline. The coils of the two end-cap

toroids are offset in ϕ from the coils of the barrel toroid by 22.5° to provide a more uniform magnetic field. The barrel and end-cap toroids provide central fields of 0.5 T and 1.0 T, respectively. The fields from the toroids point through the coils in the azimuthal direction. Muons leaving the interaction point have their trajectories bent in the R - z plane, allowing for precision measurements of the muon momentum by the Muon Spectrometer.

3.5 Inner Detector

The first detector subsystem that a particle leaving the IP will encounter is the Inner Detector (ID). The primary function of the ID is to measure the trajectories of charged particles. These trajectories are used to measure the transverse momenta of charged particles, reconstruct the position of primary and secondary vertices, and determine the charge of particles. The ID is also used to differentiate between types of particles with the same charge. It accomplishes these objectives with three complementary sub-detectors: the Pixel Detector, the Semiconductor Tracker (SCT), and the Transition Radiation Tracker (TRT). Each of these sub-detectors is arranged in concentric layers around the beamline in the barrel region and in disks in the end-caps. Figure 3.8 shows a cutaway diagram of the ID [38]. The ID has a diameter of 2.1 m and spans 6.2 m in length, providing precision tracking coverage up to $|\eta| < 2.5$.

3.5.1 Pixel Detectors

Proton bunches from the LHC cross and interact in ATLAS every 25 ns, with each bunch crossing producing thousands of particles on average. Due to the extremely high particle flux near the interaction point, a high degree of granularity is required to accurately distinguish between the individual tracks of charged particles and reconstruct them with the resolution needed to perform precision physics. The innermost component of the ID, the silicon Pixel Detector, is designed to provide precision tracking in such dense environments. The barrel region of the Pixel Detector is composed of four concentric layers of silicon pixels arranged around the beampipe between $R = 33$ mm and $R = 123$ mm. Each end-cap consists of three disks of silicon pixels oriented transverse to the beamline positioned at $|z| = 495$ mm, $|z| = 580$ mm, and $|z| = 650$ mm. Measuring $50 \mu\text{m} \times 400 \mu\text{m}$, the silicon pixels are finely segmented in $\phi - z$ and $\phi - R$ in the barrel and end-cap regions, respectively [38]. Charged particles




Figure 3.8: A cutaway diagram of the ATLAS Inner Detector [38].

passing through the pixels deposit charge in the form of electron-hole pairs, which can be measured by the readout electronics. With the exception of the innermost layer in the barrel region, the Pixel Detectors have an intrinsic hit resolution of $10\ \mu\text{m}$ in $R - \phi$ and a resolution of $115\ \mu\text{m}$ in z and R in the barrel and end-caps, respectively.

Between Run 1 and Run 2 of the LHC, the innermost of the four layers of barrel pixels was added just outside the LHC beampipe. This layer, referred to as the *Insertable B-Layer (IBL)*, significantly improves tracking and vertexing performance by reducing the distance between the first layer of the detector and interaction point [41]. The pixels of the IBL are more finely segmented than those of the other barrel layers with a nominal size of $50\ \mu\text{m} \times 250\ \mu\text{m}$ in $\phi - z$. The reduced size of the IBL pixels results in an improved hit resolution of $10\ \mu\text{m}$ in $R - \phi$ and of $67\ \mu\text{m}$ in z [42]. The IBL is critical for reconstructing tracks and vertices that are displaced from the hard scatter interaction and within the beampipe, such as the tracks from b-hadron decays. With the inclusion of the IBL, the Pixel Detector has more than 90 million readout channels. The geometry of the Pixel Detector is designed such that prompt particles originating from the interaction point with $|\eta| < 2.5$ will cross at least four

layers of pixels. A schematic of the ID geometry is shown in Figure 3.9, with an enlarged view of the Pixel Detector shown on the bottom left [43].

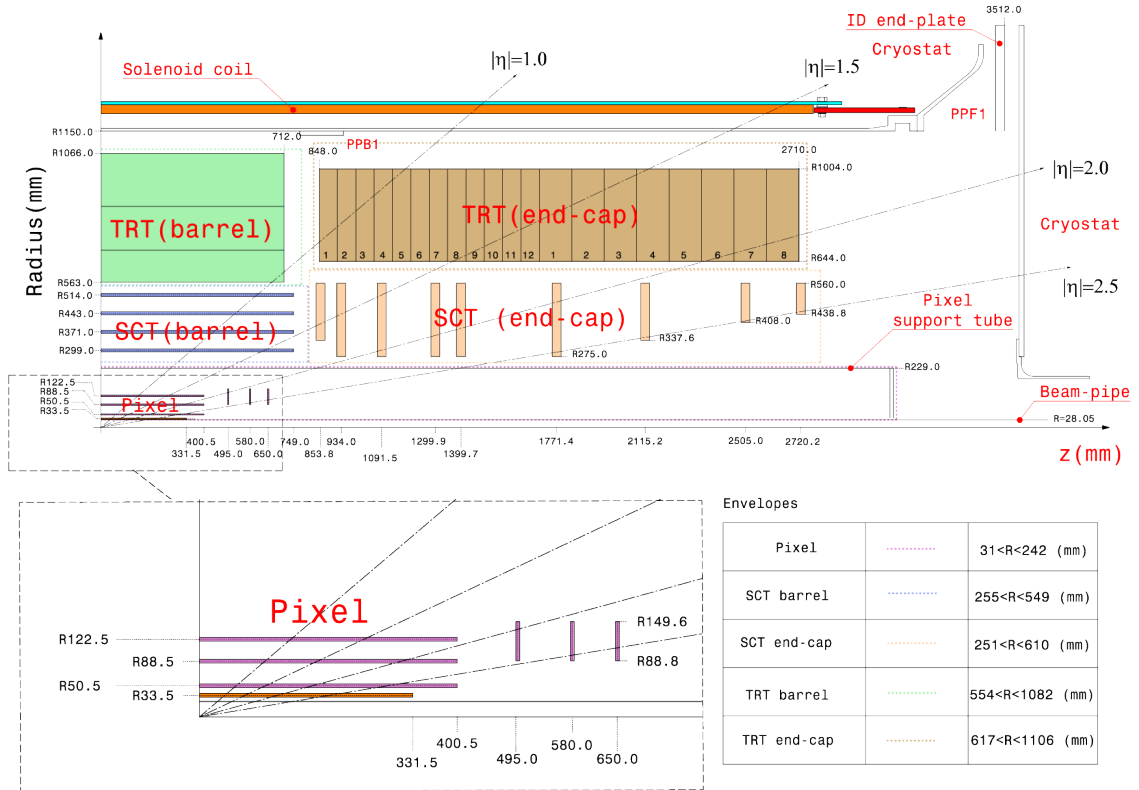


Figure 3.9: A quadrant of the ATLAS ID for Run 2 shown as an $R - z$ cross section [43]. The top diagram shows the full ID, while the bottom diagram shows a zoomed in view of the Pixel Detector.

3.5.2 Semiconductor Tracker

Like the Pixel Detector, the SCT is composed of silicon detectors arranged in concentric cylinders in the barrel and disks in the end-caps. The barrel region of the SCT consists of four layers and covers radii between 299 mm and 514 mm. Nine disks ranging from $|z| = 854$ mm to $|z| = 2721$ mm make up each end-cap [38]. Unlike the Pixel Detector, however, the SCT utilizes silicon strips rather than pixels to measure the ionization left by charged particles. Barrel SCT detectors have rectangular strips with a pitch of $80 \mu\text{m}$, while the end-cap strips are approximately trapezoidal with a mean strip pitch of $80 \mu\text{m}$. All SCT strips are between 11 cm and 13 cm in length. Each SCT module consists of two layers of silicon strips glued back-to-back. In the barrel region, each

layer contains one set of strips oriented parallel to the beamline, with the other layer rotated 40 mrad with respect to the first set. Likewise, the end-cap disks of the SCT have one set of strips that runs radially outward and another set at a slight angular offset. A single layer of silicon strips could only provide a measurement in a single coordinate, but the inclusion of a second strip layer with a small relative rotation between the layers allows the SCT to measure a second coordinate for each hit. The silicon strips of the SCT provide primary measurements with a resolution of $17\ \mu\text{m}$ in $R - \phi$, and secondary measurements of z and R with a resolution of $580\ \mu\text{m}$ in the barrel and end-caps, respectively. The geometry of the SCT, shown in Figure 3.9, is arranged such that all charged particles with $|\eta| < 2.5$ will nominally leave at least eight hits across a minimum of four SCT layers.

3.5.3 Transition Radiation Tracker

The outermost layer of the ID, the Transition Radiation Tracker (TRT), uses the ionization produced by charged particles traversing a gas to measure the trajectories of tracks. The TRT is composed of approximately 351,000 straw tube detectors in the barrel and end-cap regions [38]. Straws in the barrel region are aligned parallel to the beamline and are each 144 cm long, while straws in the end-cap TRT disks are 37 cm long and arranged radially in wheels. All TRT straws have a diameter of 4 mm. A gold-plated tungsten wire held at a voltage differential of 1530 V relative to the walls of the tube is strung down the center of each straw tube. These tubes are filled with a Xe/CO₂/O₂ gas mixture that can be ionized when a charged particle travels through it. The drift of the resulting ions in the tube induces a current in the wire that can be measured as a signal. Tracks leaving the IP with $|\eta| < 2.0$ will leave 36 hits in the TRT on average, each with an intrinsic resolution of $130\ \mu\text{m}$ in $R - \phi$. The high volume of hits provides nearly continuous tracking for charged particles at high radii.

Though the resolution of the individual hits is not as precise as the resolution offered by silicon detectors like the Pixel Detector and SCT, straw tube detectors offer two distinct advantages over silicon detectors. The first advantage is straw tubes are a significantly more cost-effective way to provide tracking over a large volume. The ATLAS TRT covers radii between 554 mm and 1082 mm. The same volume would have been extremely expensive to instrument with silicon detectors. The second key advantage over silicon detectors is that highly relativistic charged particles transitioning the boundary between two physical media emit transition radiation. The

amount of transition radiation emitted by a particle is dependent on the Lorentz factor, $\gamma = E/mc^2$. Particles with the same energy but different masses will thus leave different amounts of transition radiation in the TRT. This information is used to differentiate between particles with the same charge, such as electrons and pions. A diagram showing a cross section of the entire ATLAS ID is shown in Figure 3.10 [44].




Figure 3.10: A cross section diagram of the barrel region of the ID [44].

3.6 Calorimetry

Residing outside the ID and the central solenoid magnet is the ATLAS calorimetry system, which is designed to absorb and measure the energy of electrons, photons and hadrons for particles with $|\eta| < 4.9$. They accomplish this by initiating *electromagnetic*

and *hadronic showers*, which are cascades of secondary particles initiated by energetic particles interacting with dense material. The signals produced by the calorimetry system are additionally required to be extremely fast so that they can be used in trigger decisions (described in more detail in Section 3.8). Electromagnetic and hadronic showers differ significantly in their properties and development, so ATLAS employs two separate sub-detectors to accurately measure them: the Electromagnetic Calorimeter (ECAL), which measures the energy of electrons and photons, and the Hadronic Calorimeter (HCAL), which measures the energy of strongly interacting hadrons. A cutaway diagram of the calorimetry system is shown in Figure 3.11 [45].




Figure 3.11: A cutaway diagram of the ATLAS calorimeter system [45].

3.6.1 Electromagnetic Calorimeter

Electromagnetic showers occur when electrons or photons interact with material. In the case of electrons, this typically comes in the form of *bremstrahlung*, where photons are emitted due to the acceleration of a charged particle. Electrons passing through the dense detector material scatter off the atomic nuclei and emit *bremstrahlung*

radiation. Photons primarily interact with matter via pair production, where the photon interacts with an atomic nucleus and produces an electron-positron pair. These processes produce a cascade of particles with decreasing energy as they repeatedly occur, only stopping when the energy of the cascading photons falls below the pair production threshold ($E_\gamma < 2m_e$) and when electrons begin to lose their energy primarily through processes other than bremsstrahlung. An electromagnetic shower can be characterized by the radiation length of the medium, X_0 , which is the average distance that an electron travels in a medium before it has $1/e$ of its initial energy remaining.

The ECAL is a *sampling calorimeter*, meaning that it measures a particle's energy with a system of alternating layers of dense absorber material and sensitive active material [38]. Absorber layers are well-suited to initiating showers but cannot be used to measure the shower products, while active materials can produce signals from deposited energy but are not efficient for initiating showers. Interleaving layers of absorber and active materials allows for the showers initiated in the absorber material to be measured in the active material. A drawback of sampling calorimeters compared to *homogeneous* calorimeters (which consist entirely of active material) is that some energy from the shower is deposited in the absorber material and can't be measured. The net energy deposited by the shower cannot be used directly and must instead be estimated based on the energy deposited in the active layers.

The ATLAS ECAL employs an accordion-shaped geometry with alternating layers of lead absorber and scintillating liquid argon (LAr). The accordion geometry provides uniform coverage in ϕ with no gaps. Additionally, the placement of electrodes on the lead accordion allows for fast extraction of the signal for use in the trigger. The ECAL granularity ranges from 0.025 to 0.1 in $\Delta\phi$ and from 0.003 to 0.1 in $\Delta\eta$. The finest granularity in the ECAL is in the region where $|\eta| < 2.5$ so that precision measurements of electrons and photons can complement measurements from the ID. Both the barrel and end-cap stations of the ECAL are segmented into three layers in depth, with different granularity in each layer. This layering allows for measurements of a shower's shape as it progresses deeper into the calorimeter. Barrel modules of the ECAL have a minimum depth of $22X_0$ and end-cap modules have a minimum depth of $24X_0$ in order to contain the majority of electromagnetic showers within the ECAL [38].

3.6.2 Hadronic Calorimeter

The ATLAS HCAL is divided into three components: the Tile Calorimeter (TileCal), the LAr Hadronic End-cap Calorimeter (HEC), and the LAr Forward Calorimeter (FCAL). The TileCal is located just outside the ECAL and covers radii between 2.28 m and 4.25 m [38]. It consists of a barrel region covering $|\eta| < 1.0$ and two extended barrel regions covering $0.8 < |\eta| < 1.7$. Similar to the ECAL, the TileCal is a sampling calorimeter. Layers of steel are used to initiate hadronic showers that are then detected in layers of plastic scintillator. The development of hadronic showers in a medium is described by the nuclear interaction length λ . The nuclear interaction length is typically much longer than the radiation length that characterizes electromagnetic showers, so the HCAL is required to have greater depth than the ECAL in order to fully contain the showers. At $\eta = 0$, the TileCal has a depth of 7.4λ . The HECs use alternating layers of copper and scintillating LAr to provide hadronic energy measurements for particles with $1.5 < |\eta| < 3.2$. Each HEC has 40 LAr readout gaps divided between two wheels and a depth of approximately 10λ .

The last component of the HCAL is the FCAL, which covers the forward region of the detector ($3.1 < |\eta| < 4.9$) near the beamline. In addition to measuring the energy of forward jets, the 10 interaction lengths of material in the FCAL provide shielding for the forward regions of the Muon Spectrometer. The FCAL consists of three modules, each of which consists of wheels of absorber material oriented transverse to the beamline. Running through each wheel is a set of tubes running parallel to the beamline. Each tube contains an anode rod held at high voltage relative to the walls of the tube separated by a gap of LAr. Particles passing through the LAr produce ionization which is measured as a signal. The first of the three modules is optimized for performing electromagnetic calorimetry and the outer two modules are for hadronic measurements. The electromagnetic module uses copper as the absorber material while the hadronic modules use tungsten. A summary of the materials, pseudorapidity coverage, and depth of each component of the ATLAS calorimetry system is shown in Table 3.2.

3.7 Muon Spectrometer

Because muons are much more massive than electrons ($m_\mu/m_e \approx 206$), muons are far less likely to emit bremsstrahlung than electrons. A result of this is that muons are

Calorimeter	Absorber Material	Active Material	$ \eta $ Coverage	Depth
ECAL Barrel	Lead	LAr	$ \eta < 1.475$	$22X_0$
ECAL End-cap	Lead	LAr	$1.374 < \eta < 3.2$	$24X_0$
TileCal Barrel	Steel	Scintillator	$ \eta < 1.0$	7.4λ
TileCal Extended Barrel	Steel	Scintillator	$0.8 < \eta < 1.7$	7.4λ
HCAL End-cap	Copper	LAr	$1.5 < \eta < 3.2$	10λ
FCAL	Copper/Tungsten	LAr	$3.1 < \eta < 4.9$	10λ

Table 3.2: Materials, pseudorapidity coverage, and depth of each component of the ATLAS calorimetry system.

unlikely to initiate electromagnetic showers and thus pass through the calorimeters with minimal energy loss. The outermost layer of ATLAS is the Muon Spectrometer (MS), which is designed specifically to reconstruct muon trajectories over long distances. Muon trajectories are bent as they pass through the MS in the $R - z$ plane by the magnetic field from the toroid magnets [38].

Four different detector technologies were utilized by the MS in Runs 1 and 2. Precision tracking for muons with $|\eta| < 2.7$ is performed with Monitored Drift Tubes (MDTs) and Cathode Strip Chambers (CSCs), while Resistive Plate Chambers (RPCs) and Thin Gap Chambers (TGCs) with short signal development times are used to make trigger decisions for muons with $|\eta| < 2.4$. The trigger chambers additionally provide secondary coordinate measurements to complement the primary coordinate measurements from the MDTs. Like the ID and calorimetry systems, the MS is divided into a barrel and two end-cap regions. Three concentric cylindrical shells surrounding the beamline at $R \approx 5$ m, 7.5 m and 10 m make up the barrel region of the MS [38]. The end-caps each have three wheel stations known as the Small Wheel, the Big Wheel and the Outer Wheel, which are located at $|z| \approx 7.4$ m, 14 m, and 21.5 m respectively. An additional ring of MDTs was added to the end-caps between Run 1 and Run 2 of the LHC to provide coverage where it was previously possible for muons to only pass through two MS stations. A diagram of the MS is shown in Figure 3.12. The large distances between MS stations and the strong magnetic field provided by the toroids make it possible to extract a precise measurement of the muon momentum from the sagitta of the muon trajectory.




Figure 3.12: A cutaway diagram of the ATLAS Muon Spectrometer [38].

3.7.1 Precision Tracking Chambers

The majority of the precision tracking in the MS is provided by the MDTs, which are present in all the barrel and end-cap stations. Each MDT is a 30 mm diameter aluminum tube with a gold-plated tungsten-rhenium cathode wire in the center. The cathode wire is held at a high voltage relative to the walls of the tube, and an Ar/CO₂ gas mixture fills the tube. Muons passing through an MDT ionize the gas, and the resulting electrons drift to the wire and produce a signal. MDTs in the barrel and end-cap regions are both oriented in the ϕ -direction in order to provide a measurement resolution of 80 μm in the bending coordinate (z for the barrel and R for the end-caps). The barrel and end-cap regions each contain 20 layers of MDTs across the three detector stations. One drawback of MDTs is that the time between the ionization of the gas and the measurement of the signal can be as long as 700 ns, making the MDTs much slower than other technologies used by the MS. Due to the long drift time, information from the MDTs is not used in trigger decisions.

The forward region of the Small Wheels ($2.0 < |\eta| < 2.7$) receives the highest particle flux in the MS, and is thus the only region where precision tracking is handled by another detector technology. Cathode Strip Chambers (CSCs) are a type of multiwire proportional chamber, consisting of two cathode planes segmented into strips with a layer of anode wires strung between them [38]. An Ar/CO₂ gas mixture fills the gap between the two cathode planes. The wires of a CSC are oriented radially relative to the beamline. One set of cathode strips is oriented parallel to the wires and the other is perpendicular so that the chambers can provide measurements in two dimensions. Each CSC has four layers, which are capable of measuring the bending coordinate to a resolution of 40 μm and the transverse coordinate to a resolution of 5 mm. The time resolution of CSCs and their ability to operate safely in high rate environments made them an ideal choice for the forward region of the Small Wheels [38].

3.7.2 Trigger Chambers

In addition to the precision tracking chambers, the MS also employs two other detector technologies to provide fast measurements of muons to the trigger system. Information from Resistive Plate Chambers (RPCs) is used to make trigger decisions for muons passing through the barrel region of the MS. Two parallel resistive plates are separated by insulating spacers to form the gas gap of an RPC. Signals are induced on readout strips, with the strips of one plate oriented orthogonally to the strips of the other plate to provide measurements in two dimensions. When muons ionize the gas in the RPC, a uniform electric field of 4.9 kV/mm accelerates the free electrons, creating an avalanche that is detected on the strips. The RPCs have a typical spatial resolution of 1 cm and time resolution of 1 ns [46]. The barrel region has three RPC stations, each of which has two layers such that muons passing through the barrel of the MS can have six measurements of η and ϕ available for analysis by the trigger.

Trigger signals in the MS end-caps are handled by another type of multiwire proportional chamber called Thin Gap Chambers (TGCs). Azimuthally oriented high voltage anode wires are strung between the cathode plates to provide a strong electric field in the gas gap and a coarse measurement of the R -coordinate of the passing muons. This measurement is used to match hits to the correct precision track from the MDTs. Radially oriented strips on one of the cathode planes are used to measure the azimuthal coordinate ϕ . Three TGC stations are situated around the Big Wheel,

with an additional TGC station located on the inside of the Small Wheel. TGCs are not included in the Outer Wheel because the magnetic field does not extend to the Outer Wheel, so the ϕ -coordinate can be extrapolated from measurements at the other end-cap stations. Signals can develop quickly due to the small distance between the cathode planes and the anode wires, leading to a timing resolution of approximately 4 ns for TGCs. This is significantly shorter than the bunch crossing rate provided by the LHC. The excellent timing resolution of the TGCs provides bunch crossing identification for muons in the end-caps, allowing them to be used in the trigger decision process. A schematic showing the geometry for a quadrant of the MS is shown in Figure 3.13.

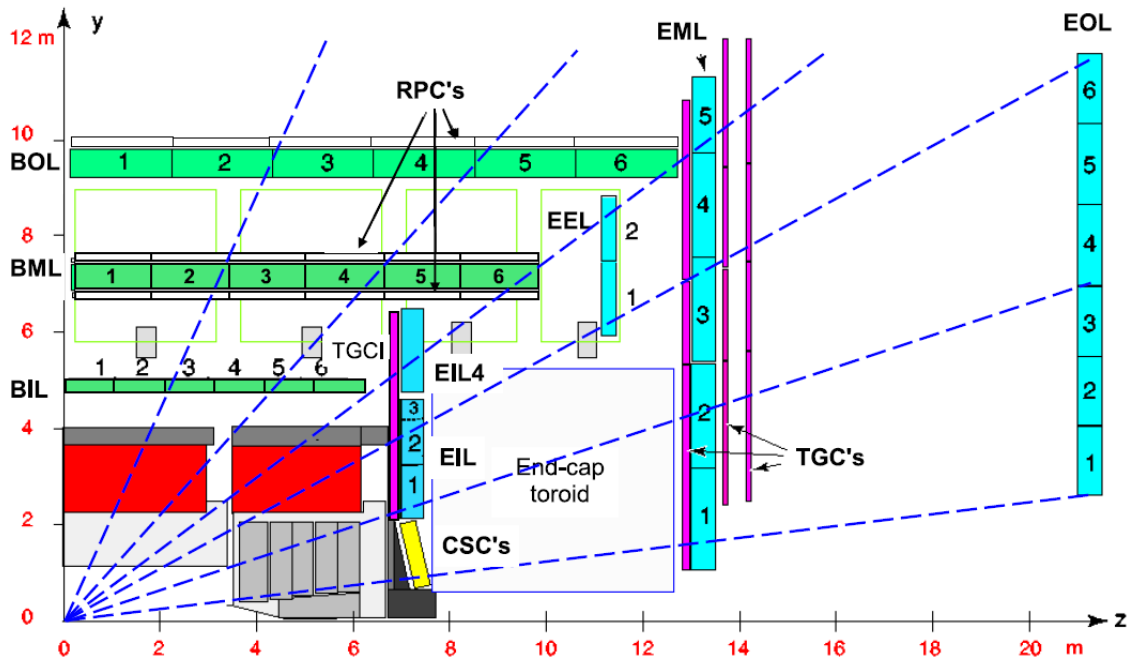


Figure 3.13: A schematic showing a quadrant of the ATLAS MS [38]. The barrel MDTs are shown in green and the end-cap MDTs are shown in blue. The CSCs are shown in yellow, TGCs are shown in purple, and the RPCs are shown as white boxes.

3.8 Trigger and Data Acquisition

Bunch crossings occur in ATLAS at a rate of 40 MHz, with an average of more than 30 pp interactions per bunch crossing [36]. The detector technologies and readout

electronics of ATLAS are unable to store every event at the LHC collision rate. Even if it was possible to record the detector output for each event, each event recorded by ATLAS contains about 1 MB of raw data output by the detector, meaning approximately 50 TB of data would be produced every second. Acquiring storage media for this volume of data in addition to the computational resources needed to reconstruct and analyze these events would be prohibitively expensive. The constraints of the readout system mean that only a small subset of events can be stored for offline analysis. Decisions on which events to keep are handled by the *trigger* system, which uses a set of simplified reconstruction algorithms to quickly identify events which may contain processes of interest to the collaboration. Such processes include events with high p_T leptons, multiple high p_T jets, or large momentum imbalances in the transverse plane.

ATLAS employed a two level trigger system during Run 2 of the LHC. The Level-1 trigger (L1) is a hardware based trigger that uses measurements from the calorimeters and the MS trigger chambers to identify candidate events for offline storage [47]. Tracking is computationally expensive, so measurements from the ID are not analyzed by the L1 trigger. Information from the calorimeters with reduced granularity is used to select candidate electrons, photons, taus, and jets with high p_T . Calorimeter measurements are also used to select events which may have large transverse momentum imbalances. Hits in the trigger chambers of the MS are used to identify muon candidates with high p_T originating from the IP. If an event passes the L1 trigger, regions of interest (ROIs) in the detector are identified for further investigation by the second trigger step. During Run 2, the L1 trigger was required to make decisions within 2.5 μ s, bringing the overall event rate down to 100 kHz.

Events passing the L1 trigger are then analyzed in greater detail by the software based high level trigger (HLT). At this stage, tracking information from the ID is incorporated into the trigger decision framework. Track reconstruction with algorithms similar to those used in offline reconstruction is carried out in the ROIs identified by the L1 trigger [48]. The full granularity of the calorimeters is also available for analysis by the HLT. Events are determined to pass the HLT if they pass one of the predefined trigger chains in the trigger menu [47]. Each trigger chain is a combination of a particular L1 trigger and a set of kinematic and object selections based on the reconstruction performed by the HLT. Trigger chains are designed to select events with certain physics signatures, such as the presence of high p_T leptons or significant missing transverse momentum. In order to prevent certain event types from consuming

too much of the total trigger bandwidth, some trigger chains have a *prescale* applied. Events passing a trigger chain with a prescale value of n have a probability of $1/n$ to be accepted by the trigger. Trigger prescales result in an effective reduction of luminosity, so analyses using prescaled triggers weight events passing those triggers by the prescale rate. The average HLT rate during a typical data collection period in Run 2 was 1.2 kHz, writing 1.2 GB/s to permanent storage [47]. Events passing the HLT are sent to the Tier-0 computing facility for offline event reconstruction [49].

Chapter 4

Event Reconstruction and Object Definition

The raw data collected by ATLAS must go through several levels of processing before it is available in a format that can be easily analyzed. The first stage of this processing is referred to as reconstruction, where the raw data is passed through a series of algorithms that use the detector output to produce the physics objects used in an analysis, such as tracks, jets, and photons. A set of standard reconstruction algorithms has been developed and optimized for the reconstruction of prompt objects. Searches for displaced vertices, however, must rely on special reconstruction algorithms because the efficiency of the standard reconstruction techniques for detecting displaced objects falls off precipitously at transverse displacements greater than a few millimeters. The $DV + E_T^{\text{miss}}$ analysis described in this dissertation utilized three special reconstruction techniques to increase sensitivity to new physics with displaced vertices. One of these special reconstruction algorithms is used to reconstruct the tracks of charged particles that do not necessarily originate from the interaction point, while the other algorithms are used to form displaced vertices from both standard and displaced tracks.

Section 4.1 will begin by outlining the concept of tracking and the techniques used to reconstruct tracks in ATLAS. The algorithms used to group tracks into primary and secondary vertices are then discussed in Section 4.2. The chapter will conclude with an overview of how other objects are reconstructed in ATLAS in Section 4.3, with special attention given to the reconstruction of jets and missing transverse energy, which are of particular importance to the search detailed in Chapters 5, 6, and 7.

4.1 Tracking

A charged particle passing through the Pixel Detector, SCT or TRT can cause ionization in the detector elements. This ionization registers a signal, referred to as a *hit*. Both the standard and large-radius tracking algorithms used by the $DV + E_T^{\text{miss}}$ analysis use the hits of Inner Detector elements to determine the trajectory taken by charged particles. These reconstructed trajectories are referred to as *tracks*.

Tracks in ATLAS are parameterized by five parameters that describe a particle's trajectory at its *perigee*, which is defined as the trajectory's point of closest approach to a reference point. These parameters and their definitions are listed in Table 4.1. An illustration of these parameters is shown in Figure 4.1. The definitions of the azimuthal angle ϕ and the polar angle θ are identical to those in Section 3.3.1. The *transverse* and *longitudinal impact parameters*, z_0 and d_0 , are defined as the transverse and longitudinal distances between a reference point and the track's perigee. The origin of the detector coordinate system serves as the reference point for track reconstruction when determining the track's perigee. The curvature of the track in the magnetic field is described by the ratio of the particle's charge to its momentum, q/p .

Parameter	Definition
d_0 - Transverse impact parameter	Distance in the transverse plane between the track's perigee and a reference point
z_0 - Longitudinal impact parameter	Longitudinal distance between the track's perigee and a reference point
ϕ - Azimuthal angle	Angle of track's momentum in the transverse plane
θ - Polar angle	Angle of track's momentum in the longitudinal plane
q/p	Charge of the track divided by the magnitude of its momentum

Table 4.1: Definitions of variables used to parameterize tracks. The reference point used to determine the track's perigee during track reconstruction is the origin of the detector coordinate system.

4.1.1 Standard Tracking

Particles originating from or near the interaction point are reconstructed by the standard tracking algorithm [51]. The first step of the algorithm clusters the charge




Figure 4.1: A schematic of the parameters used to define tracks in ATLAS [50]. The reference point for track reconstruction is the origin of the detector coordinate system.

deposited on neighboring pixels and strip sensors, and the resulting clusters are used to create three-dimensional space-points, which represent where the particle passed through the active material of the detector. Each cluster in the Pixel Detector is used to create one space-point, while SCT clusters from both sides of a strip layer are used to form a single space-point.

In the next step of the algorithm, track seeds are formed from sets of three space points, with each space-point in a seed required to be in a different layer of the Pixel Detector or SCT. Track candidate formation then proceeds through an “inside-out” approach, where track seeds are formed from space points in the Pixel Detector and SCT. From the seed, a preliminary trajectory is calculated, and the seeds are then selected or rejected based on the p_T and impact parameter of the preliminary fit. A combinatorial Kalman filter is used on surviving seeds to identify additional space-points that are compatible with the current track [52]. These compatible space-points are then added to the track.

Because all possible combinations of compatible space-points are used to make track candidates, an ambiguity solving algorithm must be run to resolve cases where space-points are shared by track candidates. A neural network is used to identify merged hit clusters resulting from the passage of multiple charged particles, and an ambiguity solver is run to limit the number of clusters shared by tracks. The ambiguity solver scores tracks using several measures of the track’s quality, such as the track χ^2 and the number of expected hits which are missing along the trajectory (referred to as *holes*), and then orders the tracks based on the resulting score. Following the ambiguity solver, all tracks are required to have no more than one shared cluster, and any cluster can be shared by no more than two tracks. Finally, tracks are extended into the TRT and compatible hits are added to the track. Tracks formed by the inside-out tracking step are referred to *Si-seeded tracks*.

Because a track’s impact parameter is defined by a point of closest approach to the beamspot, collimated tracks from displaced decays can have both a small transverse impact parameter and a displaced origin. Following the inside-out tracking step, an additional outside-in tracking step is done to reconstruct such particles. One example of this is the decay of the K_S^0 meson, which has a non-negligible lifetime and decays to $\pi^+ \pi^-$. In this tracking step, track segments from the TRT are extended back through the SCT and Pixel Detectors. Compatible hits that are not already part of a Si-seeded track are added to the TRT track segments, forming a *TRT-seeded track* [53]. All of the Si-seeded tracks and TRT-seeded tracks collectively form the collection of *standard tracks*.

4.1.2 Large-Radius Tracking

The standard tracking algorithms are able to reconstruct prompt charged particles with greater than 99% efficiency [51]. However, the tracks produced by the decays of long-lived particles may have large transverse impact parameters relative to the interaction point, resulting in a significant loss of reconstruction efficiency. To regain sensitivity to LLP signatures from potential models of new physics, an additional tracking algorithm, known as *large-radius tracking* (LRT), has been developed [54].

Large-radius tracking is performed after the standard tracking algorithm. The primary difference between LRT and the standard track reconstruction algorithm is that LRT places less stringent requirements on the track displacement from the primary vertex. Table 4.2 summarizes the difference in track selection criteria between the

standard and large-radius tracking algorithms. LRT uses hits in the silicon detectors that are not associated to a standard track to form track seeds in the same way as the inside-out step of the standard tracking algorithm. Track candidates are then constructed by using a sequential Kalman filter to extrapolate the seed trajectories. A sequential Kalman filter is chosen over the combinatorial Kalman filter used by the standard tracking algorithm in order to prevent multiple track candidates from being formed from the different combinations of space-points.

Reconstruction Step	Cut	Standard Tracking	Large-Radius Tracking
Si-Seeded Tracks	Minimum p_T	500 MeV	900 MeV
	Maximum η	2.7	5.0
	Maximum $ d_0 $	10 mm	300 mm
	Maximum $ z_0 $	250 mm	1500 mm
Clustering	Min. Si hits, not shared	6	5
	Max. # of shared hits	1	2
TRT-Seeded Tracks	Min. p_T	1000 MeV	-
	Max. $ d_0 $	100 mm	-

Table 4.2: Differences between track reconstruction criteria for standard and large-radius tracking [54].

The ambiguity solving step of the standard tracking algorithm is run on the LRT track candidates with the loosened criteria shown in Table 4.2, and the tracks which pass the ambiguity solver are extended into the TRT, where compatible TRT hits are added to the tracks. Because the TRT provides coverage for $|\eta| < 2.0$, it is not required for large-radius tracks to have a TRT extension. The LRT tracks are then merged with the standard track collection to form the final track collection for the event.

The track reconstruction efficiency for tracks originating from the decays of LLPs greatly increases when the standard tracking algorithm is complemented with LRT. The track reconstruction efficiency for charged hadrons originating from the decay of an LLP as a function of the production radius can be seen in Figure 4.2. At production radii ($r_{\text{prod}}^2 = x^2 + y^2$) greater than 50 mm, the combined efficiency can be greater than the standard efficiency by a factor of 5 or more, and significant gains in efficiency are seen for all values of r_{prod} greater than 10 mm. The combined tracking efficiency exhibits an anti-correlated dependency on the production radius, with efficiencies below 90% for production radii greater than 50 mm. The efficiency falls to approximately 35% at $R_{xy} = 300$ mm, which is the maximum allowed displacement.




Figure 4.2: Track reconstruction efficiency as a function of the production radius for displaced charged hadrons produced in the decay of long-lived BSM particles [54]. The red points represent the standard tracking efficiency, the blue points represent the LRT efficiency, and the black points show the sum of the standard and LRT efficiencies.

4.2 Vertexing

Tracks from charged particles which share a common point of origin can be clustered together to form a *vertex*. A *primary vertex* is defined as the point in space where an interaction between two protons has occurred. After the standard track reconstruction has been completed, primary vertices are reconstructed from those tracks in order to identify the location of the hard scatter interaction [55]. In order to be considered for primary vertex seeding, standard tracks must satisfy several quality requirements. In particular, strict hit multiplicity requirements are enforced in order to reduce contamination from poorly reconstructed tracks or tracks that originate from secondary interactions. The number of hits associated to a track in a given detector system is denoted as N_{PIX} , N_{SCT} and N_{TRT} for the Pixel Detectors, SCT and TRT, respectively. Because pp interactions occur inside the beampipe, at least one hit is required in either the first (IBL) or second (B-Layer) layer of the Pixel Detector for a track to be associated to a primary vertex. The number of silicon hits, N_{Si} , is the sum of N_{PIX} and N_{SCT} . Tracks are also selected based on the number of holes associated to the track, N^{holes} . The track requirements for primary vertex construction are summarized in Table 4.3.

Requirement	Value
p_{T}	$> 400 \text{ MeV}$
$ \eta $	< 2.5
$ d_0 $	$< 4 \text{ mm}$
$\sigma(d_0)$	$< 5 \text{ mm}$
$\sigma(z_0)$	$< 10 \text{ mm}$
N_{Si}	≥ 9 if $ \eta \leq 1.65$ ≥ 11 if $ \eta > 1.65$
$N_{\text{IBL}} + N_{\text{B-Layer}}$	≥ 1
# of shared modules	≤ 1 shared pixel hit or ≤ 2 shared SCT hits
$N_{\text{PIX}}^{\text{holes}}$	0
$N_{\text{SCT}}^{\text{holes}}$	≤ 1

Table 4.3: Requirements for standard tracks to be considered for primary vertex construction [55]. The uncertainties on the transverse (d_0) and longitudinal (z_0) impact parameters are denoted $\sigma(d_0)$ and $\sigma(z_0)$.

Vertex finding begins with the selection of a seed position for the vertex. The seed

position in the transverse plane is determined from the position of the beamspot, and the longitudinal seed position is taken as the mode of the track z -coordinates at their point of closest approach to the center of the beamspot. Tracks are then added to the vertex and the vertex position is refit. With each iteration, less compatible tracks are down-weighted such that they do not contribute significantly to the vertex fit. After the vertex position has been fit, tracks deemed to be incompatible with the vertex are removed from it and returned to the track collection to be used in the identification of additional primary vertices. This procedure is continued until all selected tracks have been associated to a primary vertex or no additional vertices can be constructed from the remaining tracks. The set of candidate primary vertices is composed of all vertices with at least two associated tracks. From these candidate vertices, the primary vertex (PV) is defined as the vertex with the highest sum of the squared transverse momenta of associated tracks, Σp_{T}^2 ¹. All other primary vertices are identified as pileup vertices. This selection process has proven effective at identifying the hard-scatter interaction in a given bunch crossing [55].

4.2.1 Secondary Vertexing

The primary vertexing algorithm described in Section 4.2 is designed specifically to reject vertices from displaced decays. In order to recover sensitivity to models of new physics with long-lived particles, dedicated algorithms must be used to reconstruct vertices with significant displacement from the primary vertex. Two such algorithms are utilized by the analysis described in this dissertation. The first algorithm, known as Vertex Secondary Inclusive (VrtSecInclusive, or VSI), is used in the reconstruction of point-like vertices in the Inner Detector. The second algorithm is used to reconstruct dispersed vertices from the decays of LLPs involving heavy-flavor hadrons, which have non-negligible lifetimes of their own. This algorithm, referred to as Fuzzy Vertexing (FV), is used for two of the three signal regions of the $DV + E_{\text{T}}^{\text{miss}}$ analysis. However, these regions are not the focus of the author’s work, and the Fuzzy Vertexing algorithm will consequently receive minimal discussion in this dissertation.

Secondary Vertexing - VSI

The VSI reconstruction algorithm can be divided into four steps:

¹Confusingly, pileup vertices in ATLAS are technically referred to as primary vertices. However, *the primary vertex* is the vertex meeting these selection criteria.

1. Two-track seed finding, where pairs of compatible tracks are identified to serve as seeds for DV formation.
2. N -track vertex formation, where vertices with higher track multiplicity ($N_{\text{Tracks}}^{\text{DV}}$) are formed from the existing two-track seed vertices.
3. Vertex merging, where nearby vertices are merged together if they satisfy certain criteria.
4. Track attachment, where tracks that did not satisfy the seed track requirements are considered for addition to existing DVs.

The VSI algorithm begins by identifying pairs of compatible tracks in the combined standard and large-radius track collections from which it can form two-track seed vertices [56]. The identification of secondary vertices is computationally expensive, so strict track selection criteria on the track p_{T} , $|d_0|$, and hit multiplicities are implemented to reduce the number of possible combinations available for vertex seeding. Tracks are also explicitly required not to be associated to a primary vertex from the hard-scatter or a pileup interaction. Tracks used to seed vertex formation are referred to as *selected tracks*, and they must satisfy the following criteria:

- The track must have $p_{\text{T}} > 1 \text{ GeV}$
- If $N_{\text{PIX}} = 0$, then $N_{\text{SCT}} \geq 6$
- If $N_{\text{PIX}} < 2$, then $N_{\text{TRT}} \geq 1$
- If $p_{\text{T}} < 25 \text{ GeV}$, then $N_{\text{SCT}} \geq 7$
- If $p_{\text{T}} < 25 \text{ GeV}$ and $|\eta| < 1.7$, then $N_{\text{TRT}} \geq 20$

Two-track seed vertices are then reconstructed from the set of tracks satisfying the above requirements. Pairs of tracks can make a candidate vertex if they have vertex fit quality $\chi^2/N_{\text{DoF}} < 5$ and have a radial position less than that of the SCT/TRT boundary ($R_{xy} < 563 \text{ mm}$). Both tracks are further required to pass hit-pattern checks, which requires that the tracks have hits in tracker layers located at R_{xy} greater than the vertex fit position, as well as that the tracks have no hits in tracker layers at smaller R_{xy} than the vertex fit position. These pattern requirements are not enforced for vertex candidates very close to tracker layers to avoid rejecting vertices whose position was mis-measured. Disabled silicon modules are treated as if a hit occurred to further

improve reconstruction efficiency. An illustration of the hit pattern requirements is shown in Figure 4.3 [56].




Figure 4.3: An illustration showing the hit pattern requirements imposed on tracks for a vertex decaying close to the pixel B-layer. Because the vertex position is inside the B-layer, the tracks are forbidden from having hits on the IBL and *may* have hits in the B-layer. In this example, hits are *required* in Layer 1 of the Pixel Detector [56].

Once all possible two-track vertices have been formed, the algorithm forms N -track vertices with an incompatibility graph method. Tracks are represented as nodes on the graph, and tracks that are *incompatible* with each other (i.e. tracks that *do not* form a two-track vertex) are connected by edges on the graph. The set of nodes that are fully compatible with each other can be determined by removing all irrelevant nodes from the graph, leaving a graph of completely isolated nodes. These tracks are then fit as a single multi-track vertex. Figure 4.4 shows an example of forming N -track vertices with the incompatibility graph method [57]. After this step, an individual track can still be associated to multiple displaced vertices. Tracks are required to be associated to a single vertex, so the track is removed from the vertex with which it has the poorer fit.

At this step, it is possible that multiple secondary vertices have been reconstructed from the decay products of a single LLP. Instead of a single vertex with high track multiplicity and invariant mass, multiple vertices with low track multiplicities and invariant mass may have been reconstructed. This situation is referred to as a *split vertex*. Split vertices negatively impact the sensitivity of searches that utilize the vertex invariant mass and track multiplicity to discriminate from background vertices.

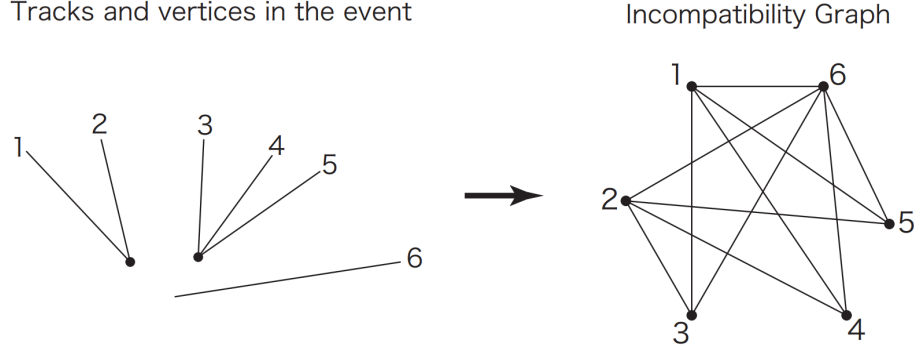


Figure 4.4: An example of identifying N -track vertices with the incompatibility graph method. The diagram on the left shows the tracks and vertices in the event, and the diagram on the right shows the incompatibility graph describing the track and vertex composition in the event [57]. Tracks are represented as nodes in the incompatibility graph, and edges represent incompatible pairs of tracks. Compatible vertices are identified by removing nodes until only isolated nodes remain.

To avoid split vertices, nearby vertices can be merged into a single vertex if they satisfy certain criteria. DVs with a distance significance $S < 10 \sigma$ are identified as candidates for merging. The distance significance is defined as

$$S = \sqrt{(\vec{v}_1 - \vec{v}_2)(\text{Cov}(\vec{v}_1) - \text{Cov}(\vec{v}_2))^{-1}(\vec{v}_1 - \vec{v}_2)^T}, \quad (4.1)$$

where \vec{v}_1 and \vec{v}_2 are the 3-D positions of the vertices and $\text{Cov}(\vec{v}_1)$ and $\text{Cov}(\vec{v}_2)$ are the covariance matrices of the vertices. Pairs of vertices satisfying the distance significance selection are merged if they meet any of the following criteria:

1. **Re-assembling:** High $N_{\text{Tracks}}^{\text{DV}}$ vertices have finer position resolution than low $N_{\text{Tracks}}^{\text{DV}}$ vertices. If all constituent tracks of the lower $N_{\text{Tracks}}^{\text{DV}}$ vertex are found to point to the high $N_{\text{Tracks}}^{\text{DV}}$ vertex after extrapolation, the DVs are merged.
2. **Suggested refitting:** The position of the lower $N_{\text{Tracks}}^{\text{DV}}$ vertex is refit using the high $N_{\text{Tracks}}^{\text{DV}}$ vertex as the starting position for the fit. If $S < 4 \sigma$ following the fit, the vertices are merged.
3. **Magnet merging:** The low $N_{\text{Tracks}}^{\text{DV}}$ vertex is refit with an additional track from the high $N_{\text{Tracks}}^{\text{DV}}$ vertex included in the fit. This is repeated for all tracks in the high $N_{\text{Tracks}}^{\text{DV}}$ vertex. If the refit position of the vertex results in $S < 4 \sigma$ between the two vertices for any of the fits, the DVs are merged.

4. **Wild merging:** A vertex is fit using all tracks from both vertices. If the vertex position is less than 4σ from the original position of the high $N_{\text{Tracks}}^{\text{DV}}$ vertex, the DVs are merged.

In a final merging step, pairs of DVs within 1 mm of each other are forced to merge. The thresholds set for the merging criteria are designed to reduce the number of split vertices while only accepting a small background contribution from the merging of unrelated vertices. Any vertices that merge through any of the above methods have their positions refit using all the constituent tracks [56]. The estimation of background from the merging of unrelated vertices is discussed in Section 6.2.1.

The final step of the algorithm is track attachment. Prior to this step in the algorithm, the reconstructed vertices consist only of tracks which satisfy the seed track criteria. However, it is possible that tracks which originated from the vertex fail the selected track selections. The inclusion of these tracks in the vertex can give a clearer kinematic picture of the vertex through their impact on the track multiplicity and invariant mass of the DV. All standard and large-radius tracks are considered, including those associated to a primary vertex. For a track to be considered for attachment to a DV, it must satisfy the following criteria:

- The track must not be associated to another DV.
- The track must have $p_T > 1 \text{ GeV}$.
- The track must have $\chi^2/N_{\text{DoF}} < 5$.
- The significance of the transverse ($|d_0|/\sigma(d_0)$) and longitudinal ($|z_0|/\sigma(z_0)$) impact parameters of the track with respect to the DV must be less than 5.
- The track must have a hit on the next outermost layer with respect to the DV position. This hit pattern check is loosened relative to the selection applied to selected tracks.

All tracks satisfying these criteria are added iteratively to each DV, beginning with the compatible DV with the highest track multiplicity. The track attachment is accepted if the vertex χ^2/N_{DoF} is less than 20 after refitting. The track attachment criteria are designed to maximize the number of attached tracks during the reconstruction phase when the track parameter extrapolation and vertex refitting information are available. The inclusivity of the algorithm is a significant strength. It

is straightforward to make analysis specific selections on the tracks at the ntuple level, but it is difficult or impossible to attach tracks to vertices at that stage.

After the track attachment step, the final vertex position is refit and the track parameters are recalculated with respect to the refitted DV position. The vertex position resolution significantly improves as the radial distance from the primary vertex decreases and as the number of tracks in the vertex increases. Typical position resolutions in the transverse plane range between $50\ \mu\text{m}$ and $300\ \mu\text{m}$ [56]. The *invariant mass* of the DV (m_{DV}) is given by the magnitude of the sum of the four-vectors of the constituent tracks. Each track is assumed to have the mass of a charged pion.

Secondary Vertexing - Fuzzy Vertexing

The VSI algorithm is highly inefficient at reconstructing decays of LLPs containing heavy-flavor quarks. Because the mesons formed by heavy-flavor quarks can have non-negligible lifetimes of their own, the tracks resulting from the original LLP decay will not necessarily appear to originate from a common point-like vertex. Instead, those tracks will appear to originate from a small volume surrounding the decay. This is the case for the Wino-Bino coannihilation model that is used as a benchmark model in the $\text{DV} + E_{\text{T}}^{\text{miss}}$ analysis, shown in Figure 2.6. A new vertexing algorithm, known as Fuzzy Vertexing (FV), has been implemented for the $\text{DV} + E_{\text{T}}^{\text{miss}}$ analysis to address the inefficiencies of the VSI algorithm in such scenarios. This search is the first ATLAS analysis to utilize FV. The first step of the algorithm is to form two-track seed vertices from pairs of tracks. Tracks are required to have p_{T} greater than 1 GeV, N_{Si} greater than or equal to two, and must not be associated to a primary or pileup vertex. Each pair of tracks is scored based on their signal-likeness by a boosted decision tree (BDT). Pairs with sufficiently high BDT scores are identified as *primary seeds*. Seed vertices near the primary seeds are then merged with the primary seeds in descending order of BDT score. Groups of seeds with common seed vertices are merged together in the final step to create a single vertex. The vertices reconstructed by the FV algorithm are referred to as fuzzy DVs or FVs. Figure 4.5 summarizes the fundamental differences between the VSI and FV algorithms.

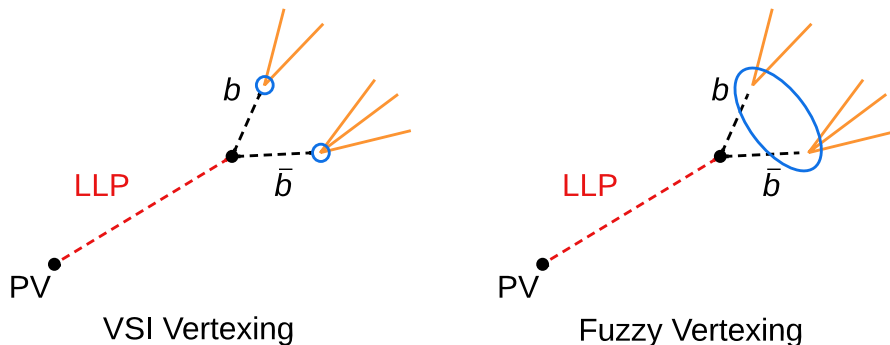


Figure 4.5: An illustration of how an LLP decaying via b -quarks would be reconstructed by the VSI and FV algorithms [58]. The b -hadron tracks are not reconstructed by the detector due to their short lifetime. Blue circles represent the vertices reconstructed by each algorithm.

4.3 Object Definitions

In addition to the tracks and vertices discussed earlier in the chapter, several other physics objects are reconstructed from the raw detector output. Following the reconstruction of these objects, parameters describing the kinematics and quality of the object can be extracted and utilized at the analysis level. This section will provide an overview of how different objects are reconstructed in ATLAS, with a focus on the objects that are most relevant to the $DV + E_T^{\text{miss}}$ analysis. A table at the end of each section will list the baseline requirements for each object used by the $DV + E_T^{\text{miss}}$ analysis. Objects passing these baseline requirements will be used in the overlap removal process described in Section 4.3.4.

4.3.1 Jets

As discussed in Section 2.2.2, quarks and gluons cannot be observed directly due to color confinement. As the distance between a strongly interacting particle and its original bound state increases, it becomes energetically favorable to create a new quark-antiquark pair from the SM fields rather than continuing to increase the binding energy between the two particles. This process, known as hadronization, produces a stream of particles travelling in the general direction of the initial particle. This stream of particles, referred to as a *jet*, can be grouped together and reconstructed with information from the calorimeters and ID. Any unbound quark or gluon in the final state of a process will result in a jet that can be measured by the detector.

EMTopo Jets

The DV + E_T^{miss} analysis described in this dissertation utilizes jets reconstructed by two algorithms. The first algorithm uses energy clusters in the calorimeters to reconstruct what will be referred to as *EMTopo jets*. The reconstruction begins by clustering calorimeter cells together into topologically connected clusters (*topo-clusters*) using a nearest-neighbor algorithm [59]. Because jets are assumed to originate from the primary vertex, an event-level correction is applied to account for the position of the primary vertex in the event.

After the cell clustering step, the topo-clusters are then reconstructed as jets by the anti- k_t jet clustering algorithm, which attempts to cluster objects into approximately conical jets with size parameter R [60]. The algorithm defines the distance between two objects i and j as

$$d_{ij} = \min(p_{T,i}^{-2}, p_{T,j}^{-2}) \frac{\Delta_{ij}^2}{R^2}, \quad (4.2)$$

and the distance between an object i and the beam as

$$d_{iB} = p_{T,i}^{-2}, \quad (4.3)$$

where $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$, and $p_{T,i}$, y_i , and ϕ_i are the transverse momentum, rapidity and azimuth of object i . For an object i , the algorithm sorts the distances to the other objects and to the beam. If the smallest distance is to another object, the two objects are clustered together, while if the smallest distance is to the beam, the object i is called a jet and removed from the list of objects. Distances are then recalculated, and the procedure is repeated until no objects remain. The anti- k_t algorithm is both infrared and collinear safe, meaning that it is insensitive to the addition of soft radiation to the event and to the splitting of particles into collinear pairs [60]. The EMTopo jets used in the DV + E_T^{miss} analysis are reconstructed with the size parameter R set to 0.4.

The four-momentum of the resulting jet is calculated as the sum of the constituent topo-cluster energy and momentum. Electromagnetic (EM) and hadronic showers in the calorimeters have sufficiently different responses and developments that an accurate measure of the deposited energy depends strongly on the type of particle initiating the shower. The type of particle initiating the shower cannot be determined with certainty, so it is assumed that all interactions are from electromagnetic processes.

This assumption biases the energy measurement of the jet, so a set of calibrations is applied to correct the energy measurement and reduce the associated uncertainty [59].

Track Jets

The second jet collection, referred to as *track jets*, constructs jets using only information from the Inner Detector. Track jets are used in the background estimation of the DV + E_T^{miss} analysis as a proxy for the track density of the event. Pileup interactions are an important contributor to background in DV searches, so the track jet reconstruction algorithm is run for all primary and pileup vertices. Standard tracks which pass the criteria in Table 4.4 are used as input to the track jet reconstruction algorithm. The selection criteria are applied to minimize the inclusion of fake tracks in the track jets. Large-radius tracks are not used in the reconstruction of track jets. Each track satisfying the selections is first associated to a primary vertex if it has longitudinal impact parameter significance $|\Delta z_0/\sigma(z_0)| < 3$ with respect to the vertex. If any tracks are not associated to a primary vertex following this initial pass, tracks are then associated to a primary vertex if they have longitudinal impact parameter $|\Delta z_0| < 0.5$ mm with respect to that vertex. Tracks that are associated to more than one vertex are removed from the vertex with lower Σp_T^2 . The tracks associated to each primary vertex are then used to reconstruct track jets with the anti- k_t algorithm with R set to 0.4, and the resulting track jets have their four-momenta calculated as the sum of the four-momenta of their associated tracks.

Requirement	Value
p_T	> 1 GeV
$ \eta $	< 2.5
$ d_0 $	< 2 mm
N_{Si}	≥ 7
$N_{\text{PIX}}^{\text{holes}}$	≤ 1
$N_{\text{Si}}^{\text{holes}}$	≤ 2
# of shared modules	≤ 1

Table 4.4: Requirements for standard tracks to be considered in the track jet reconstruction algorithm.

b-Tagging

In most cases, the jet reconstruction algorithms are not able to identify the particle that initiated the jet. However, it is possible to identify jets originating from b-quarks due to their non-negligible lifetime. Decay lengths of a few mm are typical for b-quarks produced at the LHC with significant momenta. Dedicated algorithms can use information from the Inner Detector to identify tracks with large impact parameters and associate them to jets in order to determine if a jet was initiated by a b-quark [61]. These algorithms, known as b-taggers, are utilized for the EMTopo jets used in the background estimate of the $DV + E_T^{\text{miss}}$ analysis. b-tagging can also be performed for track jets, but this is not implemented in the $DV + E_T^{\text{miss}}$ analysis.

4.3.2 Electrons and Photons

The reconstruction of electrons and photons relies on information from the Electromagnetic Calorimeter and Inner Detector. In the first step of the reconstruction algorithm, topologically connected clusters of calorimeter cells (*topo-clusters*, as discussed in Section 4.3.1) are formed in the Electromagnetic and Hadronic Calorimeters. The algorithm then attempts to match these clusters to tracks in the Inner Detector. Electrons passing through the Inner Detector are likely to lose energy through the emission of bremsstrahlung radiation, so tracks are refit to account for this energy loss. It is also possible for photons to produce an e^+e^- pair by interacting with the material of the Inner Detector. This process, known as a photon conversion, can produce a displaced vertex from the resulting e^- and e^+ tracks. These conversion vertices are also reconstructed in order to correctly identify the initial photon. Clusters which are matched to a track in the Inner Detector are identified as electrons, while unmatched clusters or clusters matched to a conversion vertex are identified as photons [62]. The energy of an electron or photon in the barrel region is estimated from a 3×7 cell area in the second Electromagnetic Calorimeter layer. A 5×5 cell area is used for the estimation in the end-cap region. In the last step of the algorithm, the electron and photon energy scales are calibrated using scale factors determined from simulated $Z \rightarrow e^+e^-$ events and $J/\psi \rightarrow e^+e^-$ events in data [63]. In the $DV + E_T^{\text{miss}}$ analysis, electrons are only used in the overlap removal process described in Section 4.3.4. Photons, however, are used for overlap removal, to veto DVs arising from photon conversions, and to provide an orthogonal trigger region for the inclusive background estimate described in Chapter 6.

4.3.3 Muons

Due to their relatively long lifetime and lack of strong interactions, muons produced at the interaction point will escape ATLAS without decaying or being absorbed by the calorimeters. As a result, it is possible for muons to leave tracks in the Inner Detector and Muon Spectrometer (MS), as well as energy deposits in the calorimeters. Muons can be reconstructed using several algorithms that take advantage of the signatures they leave in the different detector systems. However, the DV + E_T^{miss} analysis only uses two types of muons: *combined* (CB) muons and *inside-out combined* (IO) muons. Muons are used in the overlap removal process and to study the efficiency of the E_T^{miss} triggers.

The first step in the reconstruction of CB muons is the identification of muon track segments in the individual muon tracking stations [64]. Segments from different layers of the MS are then combined to form track candidates. These preliminary candidates are required to point loosely towards the interaction point. The trajectory of the muon through the magnetic field is then fit, accounting for possible interactions of the muon with the detector material. If the resulting track can be matched to a track in the Inner Detector, a combined track fit using information from both the Inner Detector and MS is performed.

In contrast, IO muons are reconstructed by extrapolating tracks from the ID to the MS. If an extrapolated track matches with three loosely-aligned MS hits, a combined fit of the ID hits, the energy loss in the calorimeters and the MS hits is performed. Because the reconstruction of IO muons does not require an independent MS track, it supplements the CB muon reconstruction efficiency in regions where the MS coverage is limited [64].

4.3.4 Overlap Removal

Because all the reconstruction algorithms discussed in Section 4.3 are run independently of one another, it's likely that the same detector information was used to reconstruct multiple particles of different types. For example, a track reconstructed in the ID could be associated to both a muon and an electron, or a calorimeter cluster could be associated to a jet and a photon. A process known as overlap removal is applied to resolve ambiguities and avoid double counting objects. This is particularly important for the calculation of E_T^{miss} , which will be discussed in Section 4.3.5. EMTopo jets, electrons, photons and muons which pass the baseline selections outlined in Table

4.5 are used as input for the overlap removal algorithm. The baseline selections can be tuned at the analysis level in order to tailor the overlap removal process towards optimal sensitivity. For the $DV + E_T^{\text{miss}}$ analysis, the following overlap removal procedure is applied:

1. **Electron-Muon:** Electrons are removed if they share a track with a muon.
2. **Electron-Jet:** Jets are removed if they are within $\Delta R < 0.2$ of an electron. Subsequently, electrons are removed if there are jets remaining within $\Delta R < 0.4$ in order to reject electrons originating from hadronic decays.
3. **Muon-Jet:** Jets which have fewer than three associated tracks are removed if they have a ghost-associated² muon or a muon within $\Delta R < 0.2$. Subsequently, muons are removed if there are jets remaining within $\Delta R < 0.4$ in order to reject muons originating from hadronic decays.
4. **Lepton-Photon:** Photons are removed if they have an electron or a muon within $\Delta R < 0.4$.
5. **Photon-Jet:** Jets are removed if they have a photon within $\Delta R < 0.4$.

4.3.5 Missing Transverse Energy

ATLAS is designed to be a hermetic detector for particles that interact through the electromagnetic and strong forces. There are particles, however, which do not interact through either of those forces, such as neutrinos, which only interact via the weak force. These weakly interacting particles can escape ATLAS completely undetected. Many models of BSM physics predict additional weakly interacting particles that would pass through ATLAS undetected. Though such *invisible* particles could not be detected directly, it is possible to infer their presence through the conservation of momentum. In the original pp interaction, the momentum of the colliding partons along the beamline cannot be known a priori due to the statistical nature of the parton momentum, governed by the parton distribution function. However, the colliding partons carry no momentum in the plane transverse to the beam. Because transverse momentum is conserved and the initial momentum in the transverse plane is known to be zero, the transverse momentum of all invisible particles in the final state can be

²The details of the ghost-association process can be found in [65, 66].

Object	Requirement	Value
EMTopo Jets	p_T $ \eta $	$> 20 \text{ GeV}$ < 4.5
Electrons	p_T $ \eta $ ID Working Point	$> 10 \text{ GeV}$ < 2.47 LooseAndBLayerLLH
Photons	p_T $ \eta $ ID Working Point	$> 25 \text{ GeV}$ < 2.37 Tight
Muons	p_T $ \eta $ ID Working Point	$> 10 \text{ GeV}$ < 2.7 Medium

Table 4.5: Baseline requirements for EMTopo jets, electrons, photons and muons. Objects passing their respective requirements will be used in the overlap removal procedure. The criteria for the identification working points are discussed in more detail in Section 5.4.2.

inferred by measuring the transverse momentum of all measured final state particles. The missing transverse momentum, also referred to as missing transverse energy, MET, or E_T^{miss} , is a powerful quantity in the search for physics beyond the Standard Model.

The missing transverse momentum vector for an event points in the direction opposite the vectorial sum of the transverse momenta of all reconstructed particles and energy deposits in the event, with equal magnitude. This can be summarized in the equation

$$\mathbf{p}_T^{\text{miss}} = - \left[\sum \mathbf{p}_T^{\text{jet}} + \sum \mathbf{p}_T^{\text{electron}} + \sum \mathbf{p}_T^{\text{muon}} + \sum \mathbf{p}_T^{\text{photon}} + \sum \mathbf{p}_T^{\text{soft}} \right], \quad (4.4)$$

where $\sum \mathbf{p}_T^{\text{soft}}$ includes all tracks and calorimeter deposits not associated with a reconstructed object [67]. The magnitude of the missing transverse momentum, denoted E_T^{miss} , is given by

$$E_T^{\text{miss}} = \sqrt{(p_x^{\text{miss}})^2 + (p_y^{\text{miss}})^2}. \quad (4.5)$$

All objects that are passed to the E_T^{miss} reconstruction algorithm are calibrated, pass the baseline selections shown in Table 4.5, and pass the overlap removal process.

Triggers based on missing momentum are unable to use the full E_T^{miss} calculation because it requires objects to be calibrated and undergo overlap removal [68]. Four algorithms were used in the HLT for Run 2 data collection:

- **cell:** E_x^{miss} and E_y^{miss} are determined from a sum over all calorimeter cells. Noise thresholds for individual cells are configured based on the average pileup of the data-taking period and the electronic noise observed in the individual cell prior to data-taking.
- **tc-lcw:** Topo-clusters, similar to those used in the jet, electron and photon reconstruction, are reconstructed and used to calculate E_x^{miss} and E_y^{miss} . Compared to the cell algorithm, this method allows clusters to be calibrated as electromagnetic or hadronic (*local cell weighting*) before calculating E_T^{miss} .
- **mht:** E_T^{miss} is calculated as the negative vector sum of the transverse momenta of all anti- k_t EMTopo jets in the event with more than 7 GeV of p_T prior to calibration. These jets have pileup subtraction and jet energy scale calibrations applied before entering the E_T^{miss} calculation.
- **pufit:** Topo-clusters in the calorimeter are combined into patches corresponding to a jet of $R = 0.4$, and a fit is performed to estimate the contribution to the cluster energy from pileup. The E_T^{miss} is then calculated as the sum of the pileup-corrected patches.

It is important to note that muons are not included in any of the E_T^{miss} calculation algorithms used in the trigger. The result of this is that muons are treated as completely invisible by E_T^{miss} triggers. Though this leads to events with fake E_T^{miss} passing the triggers, E_T^{miss} triggered events containing muons provide a valuable tool for understanding the efficiency of the E_T^{miss} trigger as the highly efficient muon triggers can be used as a reference from which to gauge the E_T^{miss} trigger efficiency. This procedure will be discussed in Chapter 6.

Chapter 5

Displaced Vertices + $E_{\text{T}}^{\text{miss}}$: Event and Object Selections

Chapters 5, 6 and 7 describe a search for long-lived particles decaying in the ATLAS Inner Detector, producing displaced vertices with high invariant mass and high track multiplicity alongside significant $E_{\text{T}}^{\text{miss}}$. This analysis, referred to as DV + $E_{\text{T}}^{\text{miss}}$, utilizes the combined 2016, 2017 and 2018 ATLAS data sets with a total integrated luminosity of 137 fb^{-1} collected at $\sqrt{s} = 13 \text{ TeV}$. This search builds on the partial Run 2 DV + $E_{\text{T}}^{\text{miss}}$ search carried out on the 2016 data set of 32.7 fb^{-1} [69]. The implementation of track attachment in the VSI algorithm, the development of the FV algorithm, improved background estimation techniques, and a data set that is more than four times larger all lead to improved sensitivity relative to the prior result. Many of these changes were implemented in the DV + Jets analysis, which searched for similar DVs in multijet final states [2]. This analysis builds further upon the groundwork laid out by the DV + Jets team.

Strict event and DV selections are applied in order to maximize the signal event yields for several benchmark models of BSM physics while remaining sensitive to other potential models that would produce DV signatures. Three *signal regions* (*SRs*) are used to efficiently select signal-like DVs while retaining a minimal amount of background:

- The **1 DV SR** requires events to contain at least one DV reconstructed by the VSI algorithm with $m_{\text{DV}} > 10 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 5$. This region was the focus of the author's work and will be detailed in Chapters 5, 6 and 7.
- The **1 FV SR** requires events to contain exactly one fuzzy vertex reconstructed

by the FV algorithm with $m_{\text{DV}} > 10 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 5$. The 1 FV SR is designed to target signals where the DV is reconstructed from the decays of heavy-flavor hadrons, such as the process shown in Figure 2.6. This region was primarily developed by other members of the analysis team and will not be discussed in detail.

- The **2 FV SR** requires events to contain at least two fuzzy vertices reconstructed by the FV algorithm with $m_{\text{DV}} > 1.5 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 4$. The 2 FV SR is designed to provide sensitivity to signals that would produce two low mass vertices from decays involving heavy-flavor hadrons. This region was primarily developed by other members of the analysis team and will not be discussed in detail.

This chapter will outline the analysis strategy, the simulated background and signal samples, and the event and DV selections used to define the 1 DV SR. Because no SM particles produce a DV with both high invariant mass and high track multiplicity, the background from this search is composed of vertices from interactions of SM particles with the detector material and algorithmic effects. These effects are poorly modeled in simulated events, so the background estimation is done in a data-driven way, detailed in Chapter 6. The results of the 1 DV SR of the analysis are presented in Chapter 7.

5.1 Motivation and Recent DV Searches

Numerous theories of physics beyond the Standard Model predict long-lived particles that decay to charged particles that could be produced at the LHC. If these particles have a lifetime between 10^{-2} ns and 10 ns , a sizable fraction of the LLPs produced in ATLAS would decay within the Inner Detector, resulting in a displaced vertex. A DV can be produced regardless of the charge of the LLP, making DV signatures a powerful tool in the search for BSM LLPs.

Several searches for displaced vertices have been carried out using ATLAS data in both Run 1 and Run 2. No dedicated DV triggers were implemented for the Run 1 or Run 2 data collection, so searches have historically been defined based on the triggering object produced in association with the DV. The first DV search using Run 2 data was the DV + $E_{\text{T}}^{\text{miss}}$ analysis performed on the 32.7 fb^{-1} 2016 data set [69]. An R -parity conserving split-supersymmetry model with pair produced long-lived gluinos, as described in Section 2.5, served as the benchmark model for the analysis [28]. A

Feynman diagram of this process is shown in Figure 2.5. The lifetime of the gluinos in this scenario is due to the mass of the virtual squarks which mediate the decay being much larger than the mass difference between the gluino and lightest neutralino.

The partial Run 2 analysis set 95% CL upper limits on the gluino pair production cross section as a function of the gluino mean proper lifetime as well as lower limits on the gluino mass as a function of lifetime [69]. Both limits were set for cases with a fixed neutralino mass of 100 GeV and a fixed mass splitting of 100 GeV between the gluino and lightest neutralino. Figure 5.1 shows the lower limits set on the gluino mass as a function of lifetime by the partial Run 2 analysis.




Figure 5.1: Lower limits on $m_{\tilde{g}}$ as a function of lifetime for (a) fixed $m_{\tilde{\chi}_1^0}$ of 100 GeV and (b) fixed $\Delta m_{\tilde{g}, \tilde{\chi}_1^0}$ of 100 GeV [69].

Following the partial Run 2 DV + E_T^{miss} analysis, two additional DV searches were done with the full Run 2 data set. The first searched for DVs produced in association with a displaced muon, which could arise from R -parity violating supersymmetry scenarios with long-lived top squarks decaying to a quark and a muon [70]. Other R -parity violating supersymmetry scenarios can result in the production of DVs alongside multiple jets. Such scenarios were targeted by the recent DV + Jets analysis [2]. A novel background estimation technique that utilized the correlation between track jets and DVs was developed for the DV + Jets analysis. A generalization of that estimation method, discussed in Chapter 6, is used as the primary background estimation method for the full Run 2 DV + E_T^{miss} analysis.

5.2 Analysis Strategy

In order to draw direct comparisons between the partial and full Run 2 DV + E_T^{miss} searches, the R -parity conserving split-supersymmetry model with pair produced long-lived gluinos is used as the benchmark for the full Run 2 DV + E_T^{miss} search. In order to keep the search as model-independent as possible, however, fine-tuning of the selections to the benchmark model is avoided. Simulated samples of the benchmark model are used to determine a set of selections that select signal efficiently across a wide range of parameter space while suppressing the number of background DVs in the 1 DV SR to approximately 1.

There are no Standard Model processes that produce high invariant mass displaced vertices in the Inner Detector, so a data-driven approach is taken to estimating the background produced from instrumental and algorithmic effects. Three main sources of background are identified:

1. **Hadronic Interactions (HI)**: Strongly interacting particles passing through the detector can interact with the nuclei in the detector material. If multiple charged particles are created via this interaction, their tracks can be reconstructed as a DV. These interactions tend to produce DVs with low invariant mass, whose decay products are generally collimated due to the boost of the incident particle with respect to the nucleus. DVs from the decays of SM LLPs such as b-hadrons are also included in this component of the background.
2. **Accidental Crossings (AX)**: DVs from decays of SM LLPs or from hadronic interactions can be crossed by an unrelated track from the primary vertex or from pileup. If the unrelated track is attached to the DV, the invariant mass and track multiplicity of the DV will increase, making it appear more signal-like.
3. **Merged Vertices (MV)**: As described in Section 4.2.1, the VSI merging algorithm attempts to merge nearby DVs based on their distance significance S . Though the merging criteria have been tuned to minimize the merging of unrelated DVs, it is possible for two DVs originating from hadronic interactions or the decay of standard model LLPs to be accidentally merged. This creates a single DV with higher invariant mass and track multiplicity.

Two approaches to the background estimation are employed. All three sources of background can be estimated inclusively by the first approach, which exploits the

correlation between DVs and the track density of the event to predict the number of background DVs. The track density estimation method is described in Section 6.1.2. Each source of background can also be estimated separately and then combined to estimate the total number of background DVs. These methods are described in Section 6.2.

5.3 Data and Simulated Samples

This analysis is performed on the combined 2016, 2017, and 2018 ATLAS data set, corresponding to an integrated luminosity of 137 fb^{-1} of $\sqrt{s} = 13 \text{ TeV}$ pp collision data. The 2015 data set of 3.2 fb^{-1} is not used for reasons that will be discussed in Section 5.3.1. A collection of simulated samples, also referred to as *Monte Carlo samples* or *MC samples*, is used to estimate signal yields, study backgrounds, and study the modeling of the $E_{\text{T}}^{\text{miss}}$ trigger efficiency.

5.3.1 Trigger and DRAW_RPVLL Filter

Performing the large-radius tracking and secondary vertexing algorithms adds significant computation time to the standard reconstruction algorithms [54]. The complexity of these algorithms makes it computationally expensive to run these special reconstruction algorithms on the entire Run 2 data set. Instead, the `DRAW_RPVLL filter` was defined by the ATLAS SUSY RPVLL (*R*-Parity Violating and Long-Lived) subgroup in order to select events of particular interest for LLP searches. After an event passes a trigger, the detector output is sent to the CERN Tier-0 computing center to undergo reconstruction. Most events are then processed with the standard reconstruction algorithms. However, if an event is identified as having passed the `DRAW_RPVLL filter`, the event instead undergoes the more computationally intensive reconstruction techniques used in LLP analyses, including LRT and VSI. During the 2016, 2017 and 2018 data-taking periods, several $E_{\text{T}}^{\text{miss}}$ triggers were included in the `DRAW_RPVLL filter`. However, the `DRAW_RPVLL filter` contained no $E_{\text{T}}^{\text{miss}}$ triggers for the 2015 data-taking period, so it was decided not to use the 2015 data set of 3.2 fb^{-1} for this analysis. This results in a 2.3% reduction in statistics.

The filter used for the DV + $E_{\text{T}}^{\text{miss}}$ analysis requires that events pass one of a selection of $E_{\text{T}}^{\text{miss}}$ triggers, shown in Table 5.1. Additionally, the event is required to have an offline $E_{\text{T}}^{\text{miss}} > 180 \text{ GeV}$ when calculated using the *LHT* algorithm (de-

noted as $\text{MET}_{\text{LocHadTopo}}$ or MET_{LHT}), which calculates the $E_{\text{T}}^{\text{miss}}$ using locally calibrated hadronic topo-clusters [71]. Events meeting these criteria are stored in the `DRAW_RPVLL` format alongside events satisfying other filters before undergoing special reconstruction, which saves the reconstructed events in the `DAOD_RPVLL` format. These files are then processed into `DAOD_SUSY15` files, which skims the files based on whether the event passed a predefined list of triggers. At this point, the data is still stored in an object-oriented framework that is not convenient for analysis. The last step of the data production process runs a custom framework which converts them to ROOT-based data structures known as *ntuples* [72]. Ntuples are ordered data sets where the data for each event is stored as a row containing n elements. Additional analysis level variables are calculated during the tuple production step. The production process from `DAOD_RPVLL` to ntuples is applied to simulated events in a similar fashion.

5.3.2 Simulated Samples

Simulated samples are used for three purposes in this analysis, the first of which is to estimate signal yields. Simulated events of the long-lived gluinos produced in the benchmark split-supersymmetry model are generated at leading order using MADGRAPH 2.6.2 [73]. The resulting events are then passed to PYTHIA 8.24 to simulate the decays, parton showering, and the underlying event [74]. Pythia uses the A14 tune and NNPDF23LO parton distribution functions [75]. No additional filters are applied to the resulting samples. These events are then passed through the GEANT4 simulation, which models the interactions of particles as they pass through ATLAS [76]. This simulation also models the collision conditions, including pileup representative of the data during each run period. The samples are then digitized and undergo the same reconstruction methods as the data, described in Chapter 4.

In order to cover a broad range of signal phase space, samples are produced with gluino masses between 400 GeV and 3 TeV. The three primary settings for the neutralino mass are $m_{\tilde{\chi}_1^0} = 100$ GeV, $\Delta m_{\tilde{g}, \tilde{\chi}_1^0} = 100$ GeV and $\Delta m_{\tilde{g}, \tilde{\chi}_1^0} = 30$ GeV. Additional samples are generated with a smaller mass splitting of $\Delta m_{\tilde{g}, \tilde{\chi}_1^0} = 10$ GeV in order to explore sensitivity in the 2 FV SR, which is beyond the scope of this dissertation. Regions of the parameter space with small mass splittings ($\Delta m_{\tilde{g}, \tilde{\chi}_1^0} \leq 100$ GeV) are referred to as *compressed* scenarios. For the above mass points, samples are generated with the gluino mean proper lifetime set between 0.01 ns and 30 ns. Gluinos with mean proper lifetime greater than 30 ns have minimal efficiency due to

Year	Period	\mathcal{L} (pb ⁻¹)	E_T^{miss} Trigger
2016	A-D3	6226.2	HLT_xe90_mht_L1XE50
	D4-	27176.0	HLT_xe110_mht_L1XE50
2017	B	5427.3	HLT_xe90_pufit_L1XE50 HLT_xe110_pufit_L1XE55
	C	2385.1	HLT_xe100_pufit_L1XE55 HLT_xe100_pufit_L1XE50 HLT_xe110_pufit_L1XE55
	D1-D5	5150.1	HLT_xe110_pufit_L1XE55 HLT_xe110_pufit_L1XE50
	D6-K	31668.1	HLT_xe110_pufit_L1XE50
2018	B	3783.5	HLT_xe110_pufit_xe70_L1XE50 HLT_xe120_pufit_L1XE50
	C-J	17479.1	HLT_xe110_pufit_xe65_L1XE50 HLT_xe110_pufit_xe70_L1XE50 HLT_xe120_pufit_L1XE50
	K-	37529.0	HLT_xe110_pufit_xe65_L1XE50 HLT_xe110_pufit_xe70_L1XE50

Table 5.1: Table of lowest unrescaled E_T^{miss} triggers by data-taking period for years 2016, 2017 and 2018. The integrated luminosity \mathcal{L} in pb⁻¹ is shown for each period. In the trigger names, the last component denotes the Level 1 trigger used in the trigger chain. For example, a trigger name ending in L1XE55 requires that the event pass the Level 1 trigger requiring 55 GeV of E_T^{miss} . The first component of the trigger name denotes the HLT trigger used in the trigger chain. These components are of the form HLT_xe{X}_{Y}, where {X} is the minimum E_T^{miss} required by algorithm {Y}. Triggers containing an additional _xe{Z} term are required to additionally have at least {Z} GeV of E_T^{miss} as determined by the cell algorithm described in Section 4.3.5.

requirements that the LLP decay within the fiducial volume ($R_{xy} < 300$ mm), and gluinos with lifetime less than 0.01 ns are covered by existing prompt searches in ATLAS.

A second set of MC samples is used to understand the background composition of DVs passing the event selections in order to more effectively discriminate against background DVs in data. Because the events must pass an E_T^{miss} trigger and have significant offline E_T^{miss} , samples of Z bosons decaying to neutrinos accompanied by additional jets are generated at leading order using SHERPA 2.2.11 [77]. These samples are divided based on the jet flavor produced in association with the Z boson - b-jets, c-jets, or light-flavor jets. Though the background for the search is estimated with data, it is instructive to look at the composition of DVs in simulated events because the truth information is available for analysis. For example, it is possible to identify DVs that come from material interactions or DVs with an accidentally crossing track using the available truth information.

The E_T^{miss} triggers used to collect data for this analysis do not reach their full efficiency with respect to the calculated offline E_T^{miss} until the offline E_T^{miss} is greater than approximately 200 GeV [68]. In order to validate the modeling of the E_T^{miss} trigger efficiency turn-on curve, a set of $Z \rightarrow \mu\bar{\mu} + \text{jets}$ samples are generated with SHERPA 2.2.2 [77]. These samples are divided based on the associated jet flavor in the same way as the $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples. The use of these samples is discussed more in Section 6.3.

5.4 Event Selections and Region Definitions

Event selections must be applied to events in data to optimize the sensitivity to DVs originating from BSM LLPs. Event selections are used to define three different types of regions in this analysis:

- **Signal Regions (SRs):** SRs are designed to have optimal sensitivity to the signal models targeted by the analysis. The selections that define these regions are chosen to have high signal reconstruction efficiency with minimal contamination from background DVs.
- **Control Regions (CRs):** CRs are regions which have a negligible number of reconstructed signal DVs compared to the expected number of DVs from background. These regions are used to develop techniques to estimate the

number of background DVs expected in the SR in an environment that will be minimally impacted by the existence of the targeted models.

- **Validation Regions (VRs):** VRs are regions used to test the performance of the background estimation technique. They should be as similar to the signal region as possible while having minimal signal contamination.

This section will describe the event-level selections used to define the regions used by the analysis, while Section 5.5 will discuss the selections that are applied to vertices to further define the regions.

5.4.1 Non-Collision Background and Tile Module Vetoes

Requiring that events pass E_T^{miss} triggers and have significant offline E_T^{miss} can create a data set that suffers from significant contamination from non-collision background (NCB). Non-collision background can originate from several processes, but the most common is tertiary beam halo, where protons from the beam escape the cleaning insertions of the accelerator [78]. These protons can then pass through the detector and leave signals in the tracker and calorimeters that are unrelated to the hard scatter. Because these protons do not originate from the PV, these signals appear as a large imbalance in the transverse energy of the event, creating significant E_T^{miss} . These events are thus likely to pass the trigger and event selections of this analysis. The particles from the beam halo do not enter the detector with a symmetric distribution in the azimuthal angle ϕ . Instead, there are strong peaks at 0 and $\pm\pi$, which lie along the horizontal x-axis of ATLAS. The result of this preference can be seen in the E_T^{miss} ϕ distribution, shown in Figure 5.2. The events shown in Figures 5.2 through 5.5 are required to pass the DRAW_RPVLL filter selections and have offline E_T^{miss} greater than 150 GeV. In Figure 5.2, the red histogram has no additional event selections applied.

To suppress NCB events, several standard cleanings have been developed [79]. The standard cleanings involve vetoing events which contain jets that fail a set of cuts defined by a working point. The working point designed for use in LLP analyses is the *SuperLooseBadLLP* working point, which vetoes *events* containing an EMTopo jet that satisfies any of the following criteria:

- ($f_{\text{max}} > 0.99$) and ($|\eta| < 2$),
- ($E_{\text{neg}} > 60 \text{ GeV}$),

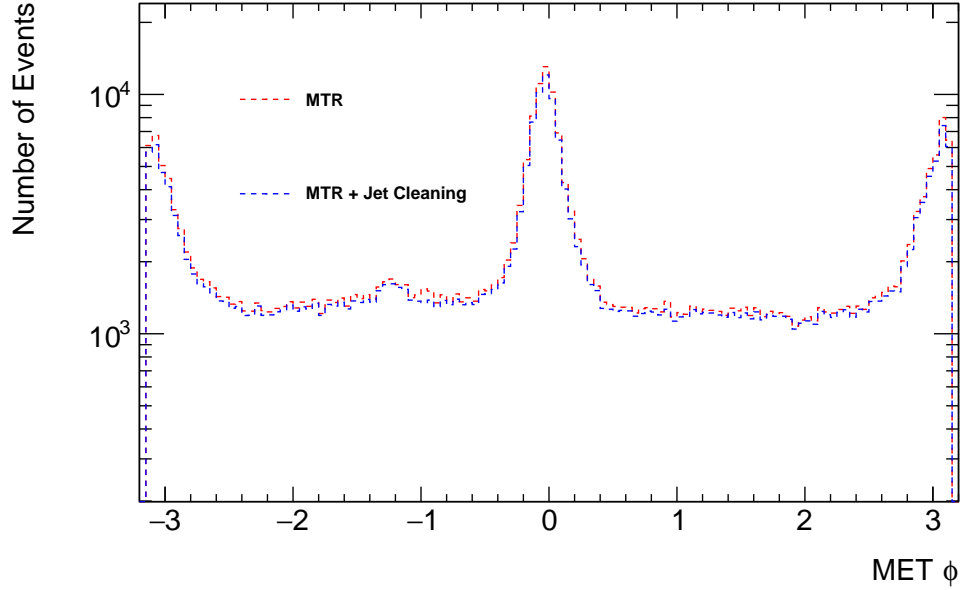


Figure 5.2: The E_T^{miss} ϕ distribution for three sample data runs (303304, 303338, 303943). Events are required to pass the DRAW_RPVLL filter selections and have offline E_T^{miss} greater than 150 GeV. The red line shows the distribution with no additional selections applied, and the blue line shows the distribution when applying the SuperLooseBadLLP jet cleaning. The jet cleaning has a negligible effect in suppressing the NCB peaks at 0 and $\pm\pi$.

- $(f_{\text{HEC}} > 0.5)$ and $(\langle Q \rangle > 0.8)$ and $(|f_Q^{\text{HEC}}| > 0.5)$,
- $(f_{\text{EM}} > 0.95)$ and $(f_Q^{\text{LAr}} > 0.8)$ and $(|\eta| < 2.8)$ and $(\langle Q \rangle > 0.8)$,

where the variables used to assess the jet quality are:

- f_{max} : the maximum energy fraction deposited in any single calorimeter layer.
- E_{neg} : the sum of the energy in all calorimeter cells with negative energy in the jet. Electronic noise or out-of-time pileup can cause calorimeter cells to have negative energy [78].
- f_{HEC} : ratio of the energy in the hadronic end-cap calorimeter to the total energy of the jet.
- $\langle Q \rangle$: the energy-squared weighted average of the pulse quality of the calorimeter cells used to build the jet.

- f_Q^{HEC} : the fraction of the energy in the HEC cells of the jet with poor signal shape quality.
- f_{EM} : the ratio of the energy in the Electromagnetic Calorimeter to the total energy of the jet.
- f_Q^{LAr} : the fraction of the energy in the liquid argon calorimeter cells of the jet with poor signal shape quality.

A small fraction of events originating from NCB can be vetoed by rejecting events which contain a jet that fails the SuperLooseBadLLP working point. The $E_T^{\text{miss}} \phi$ distribution after applying the SuperLooseBadLLP working point is shown by the blue histogram in Figure 5.2. Stricter jet cleaning working points are able to reject a greater fraction of NCB events, but these cleanings also have a significant impact on the predicted signal yield in the benchmark signal models. The partial Run 2 version of the DV + E_T^{miss} analysis included studies of several of the variables used in these cleanings in order to determine cuts that effectively suppressed NCB while retaining high signal efficiency [69]. The variables that most effectively discriminated between signal and NCB events are the f_{max} and f_{EM} of the leading jet in the event, where the leading jet is defined as the jet with the highest p_T that passed overlap removal and the SuperLooseBadLLP working point. Vetoing events with a leading jet that satisfies either $f_{\text{max}} > 0.8$ or $f_{\text{EM}} > 0.96$ effectively suppresses background without significantly affecting signal yields. Events which contain zero EMTopo jets pass the veto. Figures 5.3 (a) and 5.3 (b) show the leading jet f_{EM} and f_{max} plotted against the $E_T^{\text{miss}} \phi$, respectively. Figures 5.4 (a) and 5.4 (b) show the leading jet f_{EM} plotted against the leading jet f_{max} for sample data runs and a gluino sample, respectively. From these plots, the strong correlation between the leading jet variables and the $E_T^{\text{miss}} \phi$ distribution can be seen. A significant fraction of the events that pass the trigger and offline E_T^{miss} cuts can be removed with the selection suggested above. Figure 5.3 (a) also suggests that additional NCB events could be removed if a minimum requirement was placed on f_{EM} , but it can be seen in Figure 5.4 (b) that any minimum requirement for f_{EM} would greatly impact signal efficiency for certain regions of phase space.

In addition to the NCB peaks at 0 and $\pm\pi$ in the $E_T^{\text{miss}} \phi$ distribution, an additional peak at $\phi \approx -1.2$ is visible in Figure 5.2. During Run 2 data-taking, there were several periods where individual modules of the TileCal were not functional. Events





(a) f_{EM} vs. $E_T^{miss} \phi$ (b) f_{max} vs. $E_T^{miss} \phi$

Figure 5.3: The leading jet (a) f_{EM} and (b) f_{max} plotted against the $E_T^{miss} \phi$ for three sample data runs (303304, 303338, 303943). Events above the horizontal red line are rejected by the NCB veto.



(a) Data



(b) Simulated signal

Figure 5.4: The leading jet f_{EM} vs. f_{max} for (a) three sample data runs (303304, 303338, 303943) and (b) a gluino sample with $m_{\tilde{g}} = 2200$ GeV, $m_{\tilde{\chi}_1^0} = 100$ GeV, and $\tau = 10$ ns. Events above or to the right of the red lines are rejected by the NCB veto.

collected during these periods can have systematically mis-measured E_T^{miss} if a jet was aligned with one of the dead modules. Events containing a jet with $p_T > 50$ GeV that overlaps with a dead tile module are rejected if $\Delta\phi$ between the E_T^{miss} vector and the jet is less than 0.3. Figures 5.5 (a) and 5.5 (b) show the E_T^{miss} ϕ distributions of three data runs (303304, 303338, and 303943) for three sets of event selections:

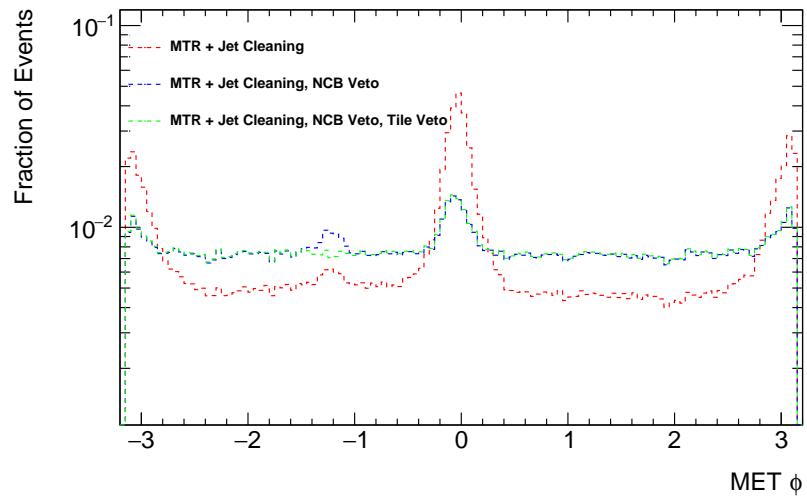
- Red histogram: events are rejected if the event contains a jet which fails the SuperLooseBadLLP cleaning criteria.
- Blue histogram: events are rejected if the event contains a jet which fails the SuperLooseBadLLP cleaning criteria, or if the event fails the NCB veto selections on f_{max} or f_{EM} .
- Green histogram: events are rejected if the event contains a jet which fails the SuperLooseBadLLP cleaning criteria, if the event fails the NCB veto selections, or if the event fails the dead tile module veto.

The plots are normalized and unnormalized, respectively. If no NCB events were present, the E_T^{miss} ϕ distribution would be approximately uniform from -2π to 2π , but this is clearly not the case without the NCB veto applied. Applying the NCB veto greatly decreases the relative amplitude of the peaks at 0 and $\pm\pi$ without significantly impacting the acceptance of events outside those peaks. The additional application of the dead tile module veto suppresses the residual peak at $\phi \approx -1.2$. After applying the NCB veto, residual peaks in the E_T^{miss} ϕ distribution remain at 0 and $\pm\pi$. NCB events are not a significant source of background DVs in CRs, VRs, or SR, so it is not critical to further suppress NCB events at the cost of signal efficiency.


5.4.2 Object Selections

All the objects defined in Chapter 4 are utilized in this analysis in some capacity. A set of additional selections are imposed to tune the object definitions for input to the overlap removal algorithm. Objects passing these selections are referred to as *baseline* objects. These selections are defined in Table 4.5, but will be summarized here.

Jets are required to have p_T greater than 20 GeV and have $|\eta|$ less than 4.5. Jets must also pass the SuperLooseBadLLP working point defined in Section 5.4.1. If an event contains a jet which fails the SuperLooseBadLLP working point, the event is discarded to avoid contamination from NCB. Track jets are required to have p_T greater than 10 GeV and at least two tracks.



(a) Normalized



(b) Unnormalized

Figure 5.5: The E_T^{miss} ϕ distribution for three sample data runs (303304, 303338, 303943). The histogram in red shows the distribution with jet cleaning applied, while the histogram in blue shows the distribution when additionally rejecting events which failed the NCB veto. The green histogram shows the distribution after additionally rejecting events which fail the dead tile module veto. The histograms in (a) are normalized to unity to show the effect on the shape of the E_T^{miss} ϕ distribution, while the histograms in (b) are unnormalized to visualize the total reduction in events from applying the vetoes.

Electrons must have p_T greater than 10 GeV and $|\eta|$ less than 2.47. Additionally, electrons must pass the ‘LooseAndBLayerLLH’ identification working point, which uses a likelihood method based on track hit multiplicities in the Inner Detector as well as on the shower shape in the calorimeter to identify an object as an electron [80]. Electrons are only used in the overlap removal and the calculation of E_T^{miss} for this analysis.

Photons are required to have p_T greater than 25 GeV and have $|\eta|$ less than 2.37. Photons must additionally pass the ‘Tight’ identification requirement [81]. The ‘Tight’ ID working point uses information from the individual layers of the electromagnetic and Hadronic Calorimeters to classify objects as photons. Photons are used in the overlap removal process and in the calculation of E_T^{miss} . Additionally, a looser set of photons is used for vetoing DVs that may have come from photon conversions. The only requirement imposed on these photons is that they have p_T greater than 60 GeV. They are not required to pass overlap removal in order to more aggressively veto DVs arising from photon conversions.

Muons are required to have p_T greater than 10 GeV and have $|\eta|$ less than 2.7. The ‘Medium’ muon identification working point imposes requirements on the hit multiplicities in the Muon Spectrometer and on the ratio of charge to momentum for the candidate muon [64]. Baseline muons are required to pass the ‘Medium’ identification criteria. Like electrons, muons are only used in the overlap removal algorithm and in the calculation of E_T^{miss} for this analysis.

5.4.3 Region Definitions

This analysis is performed in two primary regions - the E_T^{miss} -triggered region (MTR) and the orthogonal *photon-triggered region* (PTR). Each region corresponds to a set of event level selections:

- **MTR:** Events in the MTR must meet the following requirements:
 - Pass one of the E_T^{miss} triggers in Table 5.1
 - $\text{MET}_{\text{LocHadTopo}} > 180 \text{ GeV}$
 - Offline $E_T^{\text{miss}} > 150 \text{ GeV}$
 - Pass the NCB veto described in Section 5.4.1 ($f_{\text{EM}} < 0.96$ and $f_{\text{max}} < 0.8$ for the leading jet passing overlap removal in the event)

- **PTR:** Events in the PTR must:
 - Fail at least one of the MTR requirements, not including the NCB veto
 - Pass the photon trigger `HLT_g140_loose` (the event must be triggered by a photon with online $p_T > 140$ GeV)
 - Pass the NCB veto selections

Subsets of the MTR divided by the DV track multiplicity and mass are used as validation regions and the signal region for the channel of this analysis searching for VSI DVs. The PTR serves as the control region for the analysis and is used to derive the probabilities used in the inclusive background estimation method, described in Chapter 6. The PTR is similarly divided into CRs based on the DV mass and track multiplicity in the events. Tables 5.2 and 5.3 list the names of the CRs, VRs and the SR alongside the m_{DV} and N_{Tracks}^{DV} selections that define them. The coarse-binned regions defined in Table 5.2 are used as the primary regions in the inclusive background estimate, while the fine-binned regions defined in Table 5.3 are only used for diagnostic purposes. Figures 5.6 and 5.7 show a visual representation of the subdivisions in the PTR and MTR, respectively.

Region Name	Trigger Region	m_{DV} (GeV)	N_{Tracks}^{DV}
CR Low Track (CRLT)	Photon	> 10	4
VR Low Track (VRLT)	E_T^{miss}		
CR Low Mass (CRLM)	Photon	[5,10]	≥ 5
VR Low Mass (VRLM)	E_T^{miss}		
CR Extended	Photon	> 5	≥ 4
SR Extended	E_T^{miss}		
CR	Photon	> 10	≥ 5
SR	E_T^{miss}		

Table 5.2: A summary of the DV mass and track multiplicity cuts for the coarse-binned control, validation and signal regions used in the 1 DV SR inclusive background estimate. These are the primary regions used in estimating and validating the inclusive background estimate. The fine-binned regions in Table 5.3 are only used for diagnostic studies due to limited statistics.

Comparison of Events and Jets in the PTR and MTR

The inclusive background estimation method works by deriving the probability that a background DV is produced in association with a track jet or due to the track

Region Name	Trigger Region	m_{DV} (GeV)	N_{Tracks}^{DV}
CR1	Photon	[2,5]	4
VR1	$E_{\text{T}}^{\text{miss}}$		
CR2	Photon	[5,10]	4
VR2	$E_{\text{T}}^{\text{miss}}$		
CR3	Photon	[10,15]	4
VR3	$E_{\text{T}}^{\text{miss}}$		
CR4	Photon	[15,20]	4
VR4	$E_{\text{T}}^{\text{miss}}$		
CR5	Photon	[5,10]	5
VR5	$E_{\text{T}}^{\text{miss}}$		
CR6	Photon	[5,10]	6
VR6	$E_{\text{T}}^{\text{miss}}$		
CR7	Photon	[5,10]	≥ 7
VR7	$E_{\text{T}}^{\text{miss}}$		
CR8	Photon	> 20	4
VR8	$E_{\text{T}}^{\text{miss}}$		

Table 5.3: A summary of the DV mass and track multiplicity cuts for the fine-binned control and validation regions used in the 1 DV SR inclusive background estimate. These fine-binned regions are used primarily for diagnostic checks due to low statistics. The coarse-binned regions are the primary regions used for validation.




Figure 5.6: Control regions in the photon-triggered region defined by the DV mass and track multiplicity for (a) the fine-binned regions and (b) the coarse-binned regions.

density of an event. These studies can only be carried out in a region of data with negligible contamination from potential signal models. None of the signal models being considered by this analysis are produced in association with a high p_T photon. The requirements that events both fail the MTR selections and contain a high p_T photon cause the PTR to be significantly depleted of potential signal DVs. The lack of signal contamination and the general kinematic similarity to the MTR makes the PTR an excellent candidate for studying DVs arising purely from background processes.

Figure 5.8 shows a comparison of various event-level variables in the PTR (red) and MTR (blue). These include the number of track jets ($N_{\text{Track Jets}}^{\text{Event}}$) reconstructed per event (5.8a), the number of EMTopo jets ($N_{\text{EMTopo Jets}}^{\text{Event}}$) per event (5.8b), the total number of tracks associated to the primary vertex and pileup vertices ($N_{\text{Tracks}}^{\text{Event}}$) per event (5.8c), the sum of the p_T^2 for all tracks in the event associated to the primary vertex and all pileup vertices ($p_T^{2, \text{Event}}$) per event (5.8d), and the number of pileup vertices in the event (5.8e). On average, events in the PTR have more track jets than events passing the MTR selections. Conversely, the number of EMTopo jets per event is on average greater in the MTR than in the PTR. Because the MET trigger uses only calorimeter deposits for the MET calculation, it is expected that events in the MTR would have greater calorimeter activity than the PTR, and would thus have a greater number of EMTopo jets reconstructed on average.





Figure 5.7: Signal and validation regions in the E_T^{miss} -triggered region defined by the DV mass and track multiplicity for (a) the fine-binned regions and (b) the coarse-binned regions.


The MTR and PTR have similar distributions for the number of tracks associated to the primary vertex and all pileup vertices. The distribution for the MTR is slightly wider and flatter than that of the PTR, but both distributions are strongly peaked at an intermediate number of tracks (between 200 and 500). The MTR has greater fractions both of events with more than 500 tracks and events with less than 200 tracks. The $p_T^{2, \text{Event}}$ distributions are also similar, with the MTR having a slightly greater fraction of events with $p_T^{2, \text{Event}}$ greater than 650 GeV^2 . The pileup distribution shows that the MTR has a significantly higher fraction of events with more than 50 pileup vertices than the PTR.

Figure 5.9 shows a comparison of track jet and EMTopo jet variables between the PTR and MTR. The track jet p_T spectrum is comparable between the two regions, with the MTR having a slightly larger fraction of jets with p_T greater than 200 GeV. The distribution of $N_{\text{Tracks}}^{\text{jet}}$ peaks at 4 and 5 tracks in both regions, but the MTR has a greater fraction of 2- and 3-track jets and a smaller fraction of jets with 8 or more tracks.


The EMTopo jets in the PTR tend to have slightly higher p_T than those in the MTR, while the fraction of EMTopo jets that are b-tagged is approximately 50% greater in the MTR than in the PTR. The difference in the relative fraction of b-tagged




(a) Number of track jets with $p_T > 10$ GeV per event.




(b) Number of EMTopo jets with $p_T > 20$ GeV per event.



(c) Total number of tracks associated to the primary vertex and all pileup vertices per event.



(d) Sum of the p_T^2 for all tracks in the event associated to the primary vertex and all pileup vertices per event.



(e) Number of pileup vertices per event

Figure 5.8: Comparison of various event-level variables between the PTR (red) and MTR (blue). All plots are normalized, and the last bin of each histogram contains the overflow.

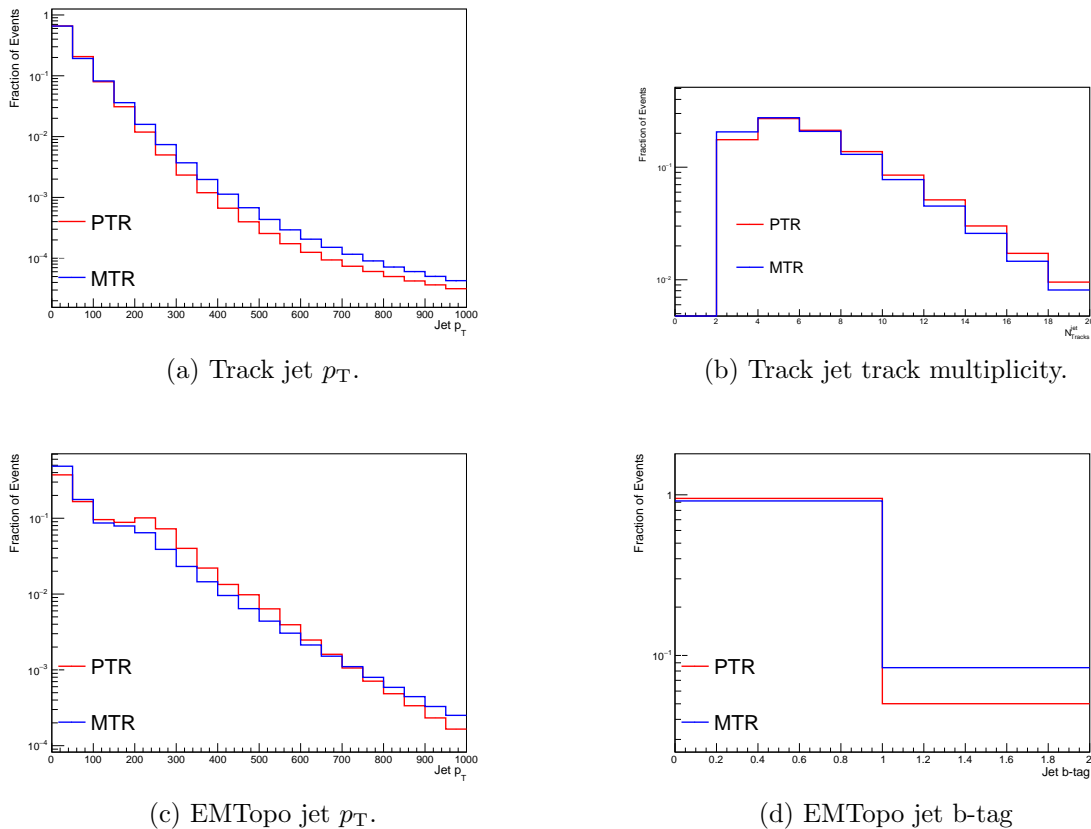


Figure 5.9: Comparison of various jet variables between the PTR (red) and MTR (blue). All plots are normalized, and the last bin of each histogram contains the overflow.

jets between the PTR and MTR is a significant factor in the estimation of background DVs. b-mesons have a non-negligible lifetime, and at LHC energies they can travel distances on the order of a few millimeters before decaying. The decays of b-mesons are a source of true SM DVs and contribute to the background of this analysis. As such, it is important for the background estimation methods to account for this difference between the two regions. The number of b-tagged EMTopo jets per event for the PTR and MTR is shown in Figure 5.10. On average, events in the MTR have twice as many b-tagged EMTopo jets as events in the PTR.




Figure 5.10: The number of b-tagged EMTopo jets in the PTR (red) and MTR (blue) in data.

5.5 Displaced Vertex Selections

Several criteria are imposed on candidate DVs to ensure the DVs are well reconstructed and to suppress background from prompt processes as well as detector and algorithmic effects. These criteria are separated into two categories - the *baseline DV* selections and the *signal DV* selections. The baseline and signal DV selections are summarized in Tables 5.4 and 5.5, respectively.

The baseline DV selections require that the DV be reconstructed within the fiducial region, defined as the volume for which $R_{xy} < 300$ mm and $|z| < 300$ mm. These constraints are determined from the DV reconstruction efficiency, which declines significantly close to $R_{xy} = 300$ mm. This is equivalent to imposing a requirement that DVs must be reconstructed before the first layer of the SCT, as shown in Figure 3.9. Additionally, all DVs are required to be separated by at least 4 mm in the transverse plane from any primary vertex in the event. This prevents primary vertices from being considered as DVs, and also suppresses background arising from heavy-flavor decays. A requirement that the DV fit χ^2/N_{DoF} be less than 5 minimizes potential contributions from fake vertices.

One potential source of secondary vertices in the Standard Model is photon

conversions into e^+e^- pairs, as discussed in Section 4.3.2. To suppress potential background from photon conversions and to make the DV composition in the PTR more comparable to the MTR, DVs are vetoed if they are reconstructed within $\Delta R_{\text{DV},\gamma}$ of 0.1 of a photon with p_T greater than 60 GeV. The photons used in this veto are not required to pass the baseline selections or overlap removal, resulting in a stricter veto. The invariant mass of a DV can sometimes be dominated by a single mis-measured track. To avoid this situation, DVs with a maximum track p_T greater than 0.95 times the scalar sum of the p_T of all tracks in the DV ($\sum_{\text{Tracks}} p_T$) are vetoed

A significant source of background DVs comes from hadronic interactions of SM particles with the material of the detector, not all of which is well modeled. These interactions can produce high track multiplicity DVs that tend to have low invariant mass. To reduce this background, DVs are rejected if they are reconstructed inside or near known material. The material veto is described in more detail in Section 5.5.1. The material veto is inverted in certain regions to validate the inclusive background estimate, as well as to develop the hadronic interactions estimate.

Fiducial volume	$R_{\text{DV}} < 300 \text{ mm}, z_{\text{DV}} < 300 \text{ mm}$
Transverse distance from all primary interactions	$R_{\text{DV-PV}} > 4 \text{ mm}$
Vertex fit quality	$\chi^2/N_{\text{DoF}} < 5$
Photon veto	$\Delta R_{\text{DV},\gamma} > 0.1$
Material veto	Outside strict material map
Max p_T veto	$\text{Max}(p_T^{\text{track}}) / \sum_{\text{Tracks}} p_T < 0.95$

Table 5.4: Summary of the baseline DV selections for the 1 DV SR.

Following the track cleaning described in Section 5.5.2, the track multiplicity ($N_{\text{Tracks}}^{\text{DV}}$), number of selected tracks ($N_{\text{Sel. Tracks}}^{\text{DV}}$), and invariant mass (m_{DV}) of the VSI vertices are recalculated. To pass the signal DV selections for the 1 DV SR, DVs are required to have at least 5 tracks passing the track cleaning selections, two of which must be selected tracks, and an invariant mass m_{DV} greater than 10 GeV. These selections significantly reduce the background in the SR, with the final number of events containing a DV in the SR being approximately one.

Invariant mass	$m_{\text{DV}} > 10 \text{ GeV}$
Track multiplicity	$N_{\text{Tracks}}^{\text{DV}} \geq 5$
Selected tracks	$N_{\text{Sel. Tracks}}^{\text{DV}} \geq 2$

Table 5.5: Summary of the signal DV selections for the 1 DV SR.

5.5.1 Material Veto

The ability to accurately identify and veto DVs arising from material interactions is critical to ensure sensitivity to potential new physics. A map of the detector material binned in R_{xy} , ϕ , and z is derived from data to determine if DVs are the result of hadronic interactions with the detector material. The map is split into an inner ($R_{xy} < 150$ mm) and outer ($R_{xy} > 150$ mm) section. The inner section corresponds to the material of the Pixel Detector, which is affixed to the LHC beampipe, and the outer section corresponds to the SCT support structure, which is connected to the ATLAS cavern. In the inner section, bins are assessed as containing material by mapping the density of displaced vertices with low invariant mass and track multiplicities. Vertices consistent with K_S^0 decays are vetoed when determining the density because they are not the result of hadronic interactions with the detector material. If the density of reconstructed DVs in a bin is higher than a predetermined threshold, the bin is flagged as containing material. In the outer section, the density of reconstructed DVs is insufficient to construct a map with the method used for the inner section. For the outer section, the material map is based on the map used by the GEANT4 detector simulation. Regions of the GEANT4 map that are not observed in data are removed from the veto and regions in data with a higher than expected density of DVs are added to the veto [70].

The veto is applied differently to data and simulated samples. The floor of the ATLAS cavern has shifted over time in relation to the beampipe. Because the SCT support structure that makes up the outer section of the material map is connected to the cavern and the inner section of the material map is connected to the beampipe, these layers have a known offset to each other when compared to the nominal ATLAS geometry. The MC simulation assumes perfect alignment of the detector, so the offsets of the inner and outer sections of the material map are corrected when applying the veto to MC samples.

Considering this, a set of four vetoes are used in the analysis: MC Loose, MC Strict, Data Loose, and Data Strict. The loose vetoes check if DVs are reconstructed in a bin containing material, accounting for the previously mentioned offsets between the inner and outer sections. This results in the rejection of approximately 42% of the fiducial volume. Figure 5.11 (a) shows the reconstructed positions of DVs in the transverse plane as seen by the ATLAS DV + Muon analysis, and Figure 5.11 (b) shows vertices that additionally fail the loose material veto [70]. The strict vetoes

impose three additional requirements on top of the loose veto criteria. First, DVs are required to not be in regions flagged as inside material by the GEANT4 simulation map in order to be conservative. The second requirement is that DVs are not reconstructed in a bin that neighbors bins flagged as inside material by the data-derived map. This is to account for the possibility that the bin contains some amount of detector material from neighboring bins, but not enough to increase the reconstructed DV density over the threshold to be flagged as inside material. The DV position uncertainty is used in the final requirement. Using the three position uncertainties of the DV, a three-dimensional ellipsoid is constructed to represent the 1σ uncertainty on the DV vertex position. DVs are then rejected if the position uncertainty ellipsoid intersects any bins flagged as containing material. These additional requirements create a conservative material veto that rejects an additional 10% of DVs when compared to the loose material veto [58].

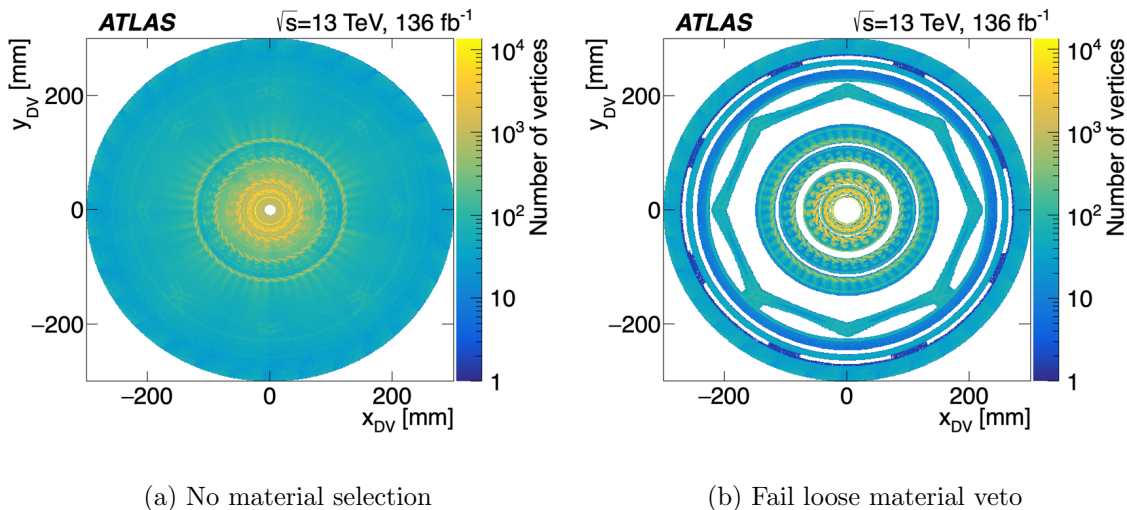


Figure 5.11: A map of reconstructed DV positions in the $x - y$ plane for DVs that pass all the baseline requirements in Table 5.4 except the photon veto and max p_T veto [70]. Vertices in (a) have no additional selections applied to them, while vertices in (b) are required to fail the loose material veto.

5.5.2 Track Cleaning

Following DV reconstruction, additional track selections are applied to the DV tracks. These selections, collectively referred to as *track cleanings*, vary depending on the

radial position of the DV the track is associated to as well as whether the track is a selected or attached track. Tracks which fail the track cleanings do not contribute towards the recalculation of $N_{\text{Tracks}}^{\text{DV}}$, $N_{\text{Sel. Tracks}}^{\text{DV}}$ and m_{DV} . Other properties of the DV, such as the DV position and goodness of fit, are not recalculated after the track cleanings. The set of track cleanings applied to VSI DVs in this analysis is summarized in Table 5.6. These selections are designed to maintain a very low background while retaining high efficiency for the models targeted by the 1 DV SR. Figure 5.12 shows the fraction of tracks remaining after each step of the track cleaning for simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ and gluino signal samples.

Name	Requirement	Track Type	Region
Transverse momentum	$p_{\text{T}} > 2 \text{ GeV}$	All	All
	$p_{\text{T}} > 3 \text{ GeV}$	Attached	Inside pixels
	$p_{\text{T}} > 4 \text{ GeV}$	Attached	Outside pixels
d_0 significance $\equiv d_0 /\sigma(d_0)$	$ d_0 /\sigma(d_0) > 10$	All	Inside beampipe
	$ d_0 /\sigma(d_0) > 15$	Attached	Inside pixels
	$ d_0 /\sigma(d_0) > 10$	Selected	Outside pixels
Angle to PV-DV vector ($\Delta\alpha$)	$\Delta\alpha > 0.02$	All	All
	$\Delta\alpha > 0.2$ if $p_{\text{T}} < 4 \text{ GeV}$	All	Outside beampipe
	$\Delta\alpha < \pi/2$	Attached	All
Backwards track veto	$\Delta\phi_{\text{PV-DV}} < 3$ if $d_0 < 1 \text{ mm}$	All	All
Upstream hit veto	No hits on layers at $R_{xy} < R_{\text{DV}}$	All	All
Hit pattern	Must pass hit pattern requirements imposed by VSI (Section 4.2.1)	All	All

Table 5.6: Selections applied to tracks in DVs reconstructed with VSI.

The majority of these cuts were originally implemented in the DV + Jets analysis [2]. Given the kinematic differences between the MTR and the region utilized by the DV + Jets analysis as well as the target signal models, it is important to confirm the efficacy of the track selections in simulated signal and background samples. In the following plots, all tracks are required to be associated to a DV which passes the fiducial volume cut, the primary vertex separation cut, and the fit quality cut. Additionally, tracks are required to pass all cleanings except those using the variables

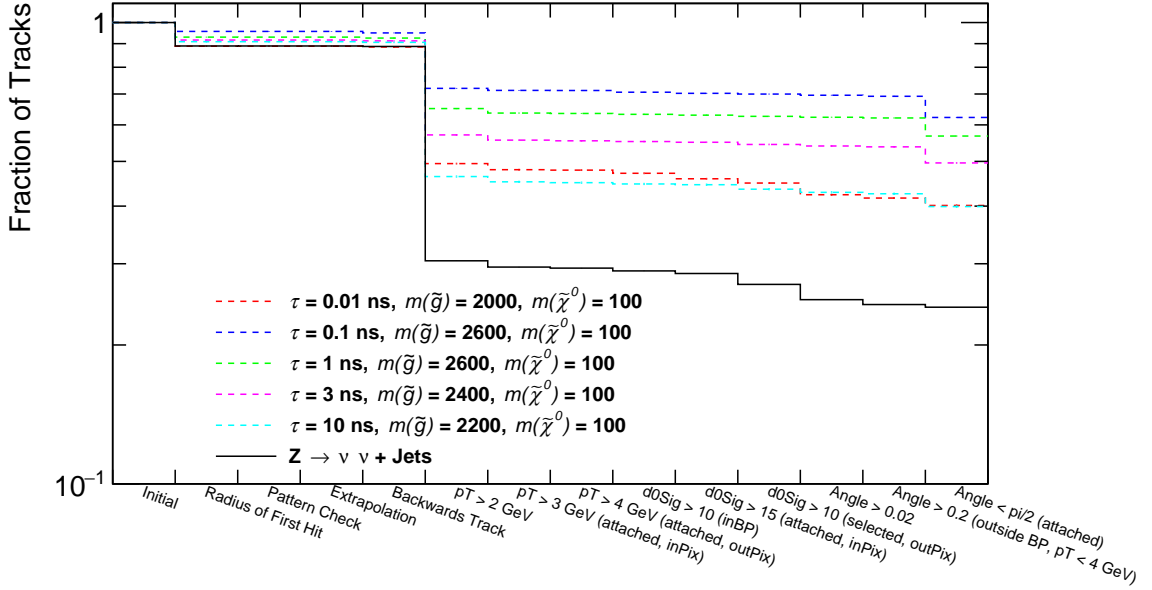



Figure 5.12: Fraction of DV tracks remaining after each track cleaning selection for simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ (black line) and several gluino signal (colored lines) samples. All tracks are required to be associated to a DV which passes the fiducial volume cut, the primary vertex separation cut, and the vertex fit quality cut.


being plotted. For example, the plots of the track p_T distribution require that tracks pass all cleanings except the p_T cleanings to be counted in the distribution. Lastly, the tracks must be in an event that passes the MTR requirements defined in Section 5.4.3.

Figures 5.13 (a), (b), and (c) show the track p_T for all tracks, attached tracks within the Pixel Detector, and associated tracks outside the Pixel Detector, respectively. The p_T distribution varies based on the radial region and whether the track is selected or attached. The rate of fake track attachment is higher in DVs at higher radii, so the p_T threshold for attached tracks increases in each successive radial region. The $|d_0|/\sigma(d_0)$ distributions for tracks inside the beampipe, attached tracks within the Pixel Detector, and selected tracks outside the Pixel Detector are shown in Figures 5.14 (a), (b), and (c), respectively. The d_0 significance cuts in Table 5.6 remove a significant fraction of background tracks while minimizing the impact on tracks originating from LLP decays.


The three-dimensional angle between the track and the PV-DV vector ($\Delta\alpha$) is a powerful discriminant against DV tracks originating from background processes. Figure 5.15 shows the distribution of $\Delta\alpha$ for tracks with different selections applied.



(a) All tracks




(b) Attached tracks within the Pixel Detector




(c) Attached tracks outside the Pixel Detector


Figure 5.13: Track p_T for DV tracks passing all other track cleanings (normalized to unity). The simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples are shown in black, while the colored lines represent simulated gluino signal samples across a range of lifetimes and mass splittings. The p_T of all tracks is shown in (a), while the p_T of attached tracks within and outside the Pixel Detector are shown in (b) and (c), respectively.



(a) All tracks within the beampipe



(b) Attached tracks within the Pixel Detector




(c) Selected tracks outside the Pixel Detector

Figure 5.14: Track $|d_0|/\sigma(d_0)$ for DV tracks passing all other track cleanings (normalized to unity). The simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples are shown in black, while the colored lines represent simulated gluino signal samples across a range of lifetimes and mass splittings. The $|d_0|/\sigma(d_0)$ of all tracks within the beampipe is shown in (a), while the $|d_0|/\sigma(d_0)$ of attached tracks within the Pixel Detector and selected tracks outside the Pixel Detector are shown in (b) and (c), respectively.

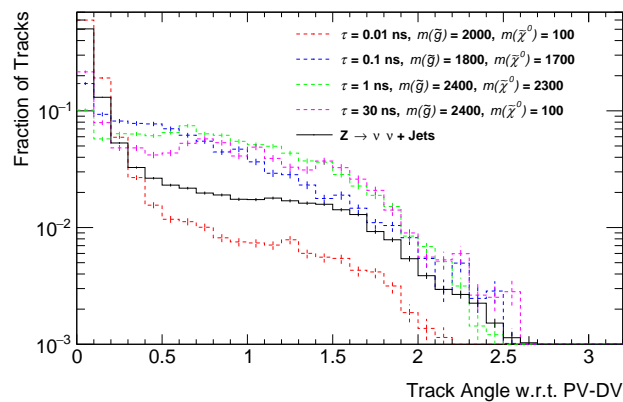
The requirement that $\Delta\alpha$ be greater than 0.02 for all tracks is imposed to reduce background from residual hadronic interactions that are not vetoed by the material map. In a hadronic interaction, the momentum of the hadron-nuclei system producing the DV is exclusively carried by the hadron. DV tracks from hadronic interactions thus tend to be highly collimated along the initial direction of the hadron. Low p_T tracks in background DVs were previously found to have small $\Delta\alpha$ compared to signal DVs, so tracks with $p_T < 4 \text{ GeV}$ are required to have $\Delta\alpha > 0.2$ [2]. Tracks at high $\Delta\alpha$ contribute significantly to the invariant mass of their associated DVs. Because tracks from signal DVs tend to travel in the same direction as their parent LLP, attached tracks are required to have $\Delta\alpha < \pi/2$ to suppress high mass background DVs arising from accidentally crossing tracks. Heavy LLPs will have less boost than lighter LLPs, leading to a more isotropic angular distribution of their decay products. Requiring $\Delta\alpha < \pi/2$ does have a significant negative impact on the signal region efficiency for heavier LLPs, but it is necessary to suppress background and maintain sensitivity. For the same reasons as above, tracks are also required to have no Inner Detector hits on tracking layers at lower R_{xy} than R_{DV} . Lastly, tracks with d_0 less than 1 mm are required to have $\Delta\phi_{PV-DV} < 3$ to prevent tracks originating from primary vertices being mistakenly included in the vertex.

5.6 Signal Yields and Efficiencies


The event, track and DV selections described in Sections 5.4 and 5.5 combine to effectively suppress background while maintaining an acceptable signal efficiency. Figures 5.16 (a) and 5.16 (b) show the selection efficiency for several R -hadron samples and the simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ background, respectively. For signals with a large mass splitting, the efficiency is primarily driven by the gluino lifetime for two reasons. DVs produced by gluinos with lifetimes less than 0.01 ns or greater than 1 ns tend to fail the PV-DV separation cut ($R_{DV-PV} > 4 \text{ mm}$) or the fiducial volume cut ($R_{DV} < 300 \text{ mm}$), respectively. DVs produced at high radii are also negatively impacted by the declining track reconstruction efficiency for tracks with significant displacement (shown in Figure 4.2). The optimal selection efficiency is therefore for gluino samples with lifetimes around 0.1 ns. The mass and track multiplicity requirements can significantly impact the selection efficiency for certain regions of the signal phase space, but those cuts are the key to ensuring that little to no SM background is present in the signal region.



(a) All tracks



(b) Tracks outside the beampipe



(c) Attached tracks

Figure 5.15: Track angle with respect to the PV-DV vector ($\Delta\alpha$) for DV tracks passing all other track cleanings (normalized to unity). The simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples are shown in black, while the colored lines represent simulated gluino signal samples across a range of lifetimes and mass splittings. The distribution of $\Delta\alpha$ is shown for (a) all tracks, (b) tracks associated to DVs outside the beampipe, and (c) attached tracks.

Figures 5.17, 5.18, and 5.19 show the yields and efficiencies as a function of the gluino lifetime and mass for a fixed neutralino mass of 100 GeV, a fixed $\Delta m_{\tilde{g},\tilde{\chi}_1^0}$ of 100 GeV, and a fixed $\Delta m_{\tilde{g},\tilde{\chi}_1^0}$ of 30 GeV, respectively. Efficiency is generally substantially higher for samples with a large value of $\Delta m_{\tilde{g},\tilde{\chi}_1^0}$, as small mass splittings have less intrinsic E_T^{miss} from the final state neutralino to satisfy the trigger and offline E_T^{miss} requirements. At intermediate lifetimes near 0.1 ns, the selection efficiency reaches approximately 40% for samples with large mass splittings. The selection efficiency can be as low as 0.01% for samples with long lifetimes or small mass splittings.






(a) R -hadron samples with $m_{\tilde{\chi}_1^0} = 100$ GeV(b) $Z \rightarrow \nu \bar{\nu}$ + jets samples

Figure 5.16: Efficiency cutflows for simulated (a) gluino R -hadron signal samples with fixed $m_{\tilde{\chi}_1^0} = 100$ GeV and (b) $Z \rightarrow \nu \bar{\nu}$ + jets samples. The signal selections are most efficient for gluinos with a lifetime between 0.1 ns and 3 ns. No simulated $Z \rightarrow \nu \bar{\nu}$ + jets events satisfy the signal region selections. The $Z \rightarrow \nu \bar{\nu}$ + jets samples are filtered by the flavor of the jets in the event.


(a) Expected number of signal events for an integrated luminosity of 137 fb^{-1} 

(b) Selection efficiencies

Figure 5.17: Predicted yields (a) and selection efficiencies (b) in the 1 DV SR for gluino samples with a fixed neutralino mass of 100 GeV. The y-axis is $m_{\tilde{g}}$ and the x-axis is the mean proper lifetime τ of the gluino. The expected number of signal events is normalized to an integrated luminosity of 137 fb^{-1} .





(a) Expected number of signal events for an integrated luminosity of 137 fb^{-1}



(b) Selection efficiencies

Figure 5.18: Predicted yields (a) and selection efficiencies (b) in the 1 DV SR for gluino samples with a gluino-neutralino mass difference of 100 GeV. The y-axis is $m_{\tilde{g}}$ and the x-axis is the mean proper lifetime τ of the gluino. The expected number of signal events is normalized to an integrated luminosity of 137 fb^{-1} .

(a) Expected number of signal events for an integrated luminosity of 137 fb^{-1} 

(b) Selection efficiencies

Figure 5.19: Predicted yields (a) and selection efficiencies (b) in the 1 DV SR for gluino samples with a gluino-neutralino mass difference of 30 GeV. The y-axis is $m_{\tilde{g}}$ and the x-axis is the mean proper lifetime τ of the gluino. The expected number of signal events is normalized to an integrated luminosity of 137 fb^{-1} .

Chapter 6

Displaced Vertices + E_T^{miss} : Background Estimation and Uncertainties

No known SM processes produce displaced vertices with both high invariant mass and high track multiplicity. Background vertices for the DV + E_T^{miss} search instead arise from a combination of detector and algorithmic effects. These effects tend to be poorly modeled by simulated samples, so a data-driven approach is taken to estimate the expected number of background events in the 1 DV SR. As described in Section 5.2, the three primary sources of background DVs are hadronic interactions (HI), merged vertices (MV), and accidentally crossed vertices (AX). A pair of complementary approaches to estimating the expected number of background DVs is used in this analysis. The primary method estimates all three background sources inclusively by exploiting the correlation between background DVs and the track density of the event. An alternative method where each source of background is estimated with a separate technique, referred to as the *combined method*, is used as a cross-check of the inclusive method. Each of the background estimation methods is validated in regions of data where the material map veto has been inverted and in sideband regions where the m_{DV} or $N_{\text{Tracks}}^{\text{DV}}$ selections have been changed. The inclusive method is described in Section 6.1, while the estimation of the individual background components is detailed in Section 6.2. Section 6.3 will conclude this chapter with a discussion of several systematic uncertainties that will be assessed on the predicted signal yields.

6.1 Inclusive Background Estimation

The underlying assumption of the inclusive background method is that the likelihood that an event contains a DV is correlated with certain properties of the event that describe its track density. This section describes three variations of the inclusive method that have been studied for use in this analysis:

1. The *jet matching method*, which uses the correlation between DVs and track jets to estimate the number of DVs in a region. This method was originally developed and implemented for the ATLAS DV + Jets analysis [2].
2. The *track density method*, which uses the correlation between DVs and the track density of the event to estimate the number of DVs in a region.
3. The *hybrid method*, which performs jet-DV matching with EMTopo jets and attributes any unmatched DVs to the track density of the event in order to derive the probability of a DV being produced in an event.

Each inclusive method begins by deriving a probability that a DV will be produced in a given event or jet in a control region orthogonal to the MTR, which contains the 1 DV SR. The photon-triggered region (PTR) described in Section 5.4.3 is chosen to serve as the control region due to its kinematic similarity to the MTR and its lack of signal contamination. The track density method described in Section 6.1.2 is the *nominal background estimation method* for the 1 DV SR. The jet matching method is discussed to provide historical context and motivation for the track density method, and the hybrid method is used as a cross-check of the track density method. Unless otherwise specified, all plots in Section 6.1 include overflow in bins bordering the top or right edge of the plot.

6.1.1 Jet Matching Method

The original iteration of the inclusive method, the jet matching method, was developed for the DV + Jets analysis, which searched for DVs in multijet final states. The method works under the assumption that the presence of DVs in an event is correlated with the track density of the event, which is itself correlated with the presence of track jets in the event. This correlation between DVs and track jets in data was investigated and confirmed in the DV + Jets analysis [2].

The jet matching method is performed in three main steps. In the first step, DVs with the target mass and track multiplicity in the PTR are matched to the closest track jet in the event, measured in terms of ΔR between the DV and the track jet. These track jets are required to have $p_T > 10$ GeV, as described in Section 4.3.1. There is no maximum $\Delta R_{\text{DV,Jet}}$ requirement for a DV to be matched to a jet.

After DVs have been matched to the nearest track jet in the event, the DVs are binned into a histogram (H_{Matched}) based on the properties of their matched track jet. The default variables used in the binning are the jet p_T and the jet track multiplicity ($N_{\text{Tracks}}^{\text{jet}}$). A histogram of all track jets in the PTR with the same binning as the DV histogram is also created at this step (H_{PTR}). The same histogram is also produced for track jets in events passing the MTR (H_{MTR}). The histogram H_{Matched} is divided by the histogram H_{PTR} to produce a histogram of the *jet-DV probability* (JDP). The JDP histogram (H_{JDP}) is parameterized as a function of the variables used to bin the histograms. The probability that a DV is produced in association with a given jet is calculated in the PTR with

$$H_{\text{JDP}} \equiv P(\text{DV} \mid \text{jet}(p_T, N_{\text{Tracks}}^{\text{jet}})) = \frac{H_{\text{Matched}}}{H_{\text{PTR}}}. \quad (6.1)$$

The track density method described in Section 6.1.2 uses an analogous method to calculate the probability that a DV is produced in a given event. An example probability calculation for the track density method is shown in Figure 6.3.

In the final step, the JDP is assigned as a weight to all track jets in events passing the MTR selections by multiplying H_{JDP} by H_{MTR} . The integral of the resulting histogram provides an estimate of the number of expected DVs in the target region. This can be expressed as

$$N_{\text{Bkg}} = \sum_{i=0}^{N_{\text{Track Jets}}} P(\text{DV} \mid \text{jet}_i), \quad (6.2)$$

where each jet passing the MTR selections is assigned a weight based on its p_T and $N_{\text{Tracks}}^{\text{jet}}$. The estimate is then given by the sum of the jet weights. It is important to emphasize that when using Equations 6.1 and 6.2 to estimate the number of DVs in a given mass and track multiplicity range, the JDP must be calculated only using DVs in the PTR that also satisfy the same mass and track multiplicity criteria as the target region. This requirement is enforced because the shape of the probability distribution with respect to the parameterization variables differs depending on the

mass and track multiplicity of the target DVs.

Due to the rarity of DVs which pass the baseline DV selections, it is unlikely to have a histogram with well-populated bins for any region except those where the DVs have low invariant mass ($m_{\text{DV}} > 5 \text{ GeV}$) or low track multiplicity ($N_{\text{Tracks}}^{\text{DV}} \leq 4$). Empty bins in a JDP histogram imply that the probability for a jet in that bin to produce a DV with the given properties is zero. This is likely not to be the case. Rather, it is the product of the extreme rarity of DVs. For most ranges of m_{DV} and $N_{\text{Tracks}}^{\text{DV}}$, the number of DVs observed in a given bin of the histograms is a Poisson process where zero is a likely outcome. This could potentially be resolved through the use of coarser bins in the histograms. However, studies using coarser bins for the histograms find that dependencies of the probability on the parameterization variables seen with finer bins are lost when using a coarser binning.

To compensate for the low number of DVs observed in the PTR, the estimate can be performed using an extended statistics method. This is done with a simple modification of Equation 6.2,

$$N_{\text{Bkg}} = f \cdot \sum_{i=0}^{N_{\text{Track Jets}}} P_{\text{Extended}}(\text{DV} | \text{jet}_i), \quad (6.3)$$

where $P_{\text{Extended}}(\text{DV} | \text{jet}_i)$ is the JDP for a loosened DV selection on m_{DV} and $N_{\text{Tracks}}^{\text{DV}}$ and f is an additional weight which takes into account the difference in DV mass and track multiplicity requirements used to calculate the JDP. This additional weight is calculated separately for each region. f is calculated for a region CR_X by dividing the number of DVs that satisfy the CR_X selections in the PTR by the number of DVs that satisfy the extended m_{DV} and $N_{\text{Tracks}}^{\text{DV}}$ selections in the PTR:

$$f_{\text{CR}_X} = \frac{N_{\text{DV}}(\text{Pass CR}_X \text{ Selection})}{N_{\text{DV}}(\text{Pass Extended Selection})}. \quad (6.4)$$

The extended selection for the 1 DV SR considers DVs with $m_{\text{DV}} > 5 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 4$. As an example, the weight f for the 1 DV SR is given by the number of DVs in the PTR with $m_{\text{DV}} > 10 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 5$ divided by the number of DVs in the PTR with $m_{\text{DV}} > 5 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 4$:

$$f_{\text{SR}} = \frac{N_{\text{DV}}(m_{\text{DV}} > 10 \text{ GeV}, N_{\text{Tracks}}^{\text{DV}} \geq 5)}{N_{\text{DV}}(m_{\text{DV}} > 5 \text{ GeV}, N_{\text{Tracks}}^{\text{DV}} \geq 4)}. \quad (6.5)$$

Shortcomings of the Jet Matching Method

Several issues arise when implementing the jet matching method into the analysis:

1. DVs are sometimes matched to track jets that are far away in ΔR , as can be seen in Figure 6.1. This was less of an issue for the DV + Jets analysis where the jet matching method was originally developed due to the multijet requirements of the final state [2]. The high track jet multiplicity of the events made it more likely that DVs would be matched to nearby jets. The lack of a maximum ΔR requirement in the track jet matching criteria allows for DVs to be matched to track jets on the other side of the detector. Background DVs are probably not related to these track jets in a meaningful way.
2. There are events containing DVs in the MTR and PTR that contain no track jets. Again, this did not occur for the DV + Jets analysis due to the final state requiring multiple jets. A minimum number of track jets was also required for events in the version of the PTR employed by that analysis. For this search, however, no such requirement is enforced, leading to the possibility of DVs going unmatched to track jets. The fraction of unmatched DVs can be non-negligible for some DV masses and track multiplicities.
3. The jet matching method significantly underestimates the number of DVs with 4 tracks and mass between 2 and 5 GeV (VR1) in the MTR. The method also significantly underestimates the number of DVs inside material for several ranges of DV mass and track multiplicity. The underestimation becomes more significant for regions with higher DV masses. The expected background predicted by the jet matching method in the outside material VRs is shown in Figure 6.2, and the observed number of DVs in the same VRs is shown in Figure 6.12 (b).

These shortcomings motivate attempts to find a more generalized method that more directly quantifies the probability of an event containing a DV based on the track density of the event, rather than using track jets as a proxy for the track density.

6.1.2 Track Density Method

To address the shortcomings of the jet matching method, two alternative methods have been developed as extensions of the original inclusive estimate. The first method, referred to as the track density method, is designed to address the problem of DVs in




Figure 6.1: ΔR between DVs and the closest track jet in the PTR. DVs are required to pass all baseline selections and have $m_{DV} > 5 \text{ GeV}$ and $N_{Tracks}^{DV} \geq 4$. 23% of DVs in the extended signal region (14 out of 62) have $\Delta R_{DV,Jet} > 0.4$.




Figure 6.2: The number of DVs predicted by the jet matching method in the signal and validation regions. The uncertainties are obtained by propagating the statistical uncertainty on the number of DVs in the PTR used to derive the JDP.

events that do not contain track jets. The assumption of the jet matching method is that the presence of background DVs is correlated with track density, and track jets are used as a proxy for the track density of the event. The track density method again assumes the correlation between track density and background DVs, but it does not match DVs to individual track jets to derive a jet-DV probability. Instead, the method derives an event-DV probability (EDP) based on variables that directly describe the track density and content of the event.

The probability that a DV is produced in an event is calculated in the PTR with


$$P(\text{DV} \mid \text{event}) = \frac{\# \text{ of DVs in events with } (x, y)}{\# \text{ of events with } (x, y)}, \quad (6.6)$$

where x and y are variables describing the track density. Such variables include the total number of tracks associated to the primary vertex and all pileup vertices ($N_{\text{Tracks}}^{\text{Event}}$), the summed p_{T}^2 of all tracks associated to the primary vertex and all pileup vertices ($p_{\text{T}}^{2, \text{Event}}$), the number of track jets reconstructed in the event ($N_{\text{Track Jets}}^{\text{Event}}$), and the number of b-tagged EMTopo jets in the event ($N_{\text{b-tag}}$). After calculating the EDP in the PTR (parameterized by the chosen track density variables), the EDP is then applied to events passing the MTR event selections in order to estimate the number of DVs in the MTR. An example of this workflow is shown in Figures 6.3 and 6.4. The variables chosen to parameterize the EDP for the final estimate are $N_{\text{Tracks}}^{\text{Event}}$ and $N_{\text{b-tag}}$. These variables are found to provide the most accurate estimate when comparing to the number of observed events in the validation regions. All results of the track density method are shown using this parameterization.

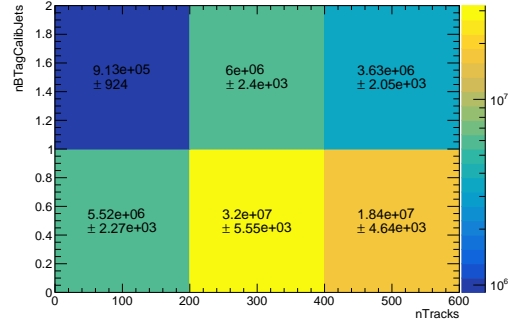
The implementation of the track density method is analogous to that of the jet matching method. The number of background DVs with a given mass and track multiplicity is estimated by assigning each event that passes the MTR selections a weight $P(\text{DV} \mid \text{event})$ equal to the EDP derived in the PTR corresponding to the given DV mass and track multiplicity. The sum of the weighted events gives the estimated number of DVs with that mass and track multiplicity:

$$N_{\text{Bkg}} = \sum_{i=0}^{N_{\text{Events}}} P(\text{DV} \mid \text{event}_i). \quad (6.7)$$


As with the jet matching method, the EDP is calculated using only DVs in the PTR that also satisfy the mass and track multiplicity requirements of the region that is



(a) The number of DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5 \text{ GeV}$ outside material in the PTR, binned by $N_{\text{b-tag}}$ and $N_{\text{Tracks}}^{\text{Event}}$.




(b) All events in the PTR binned by $N_{\text{b-tag}}$ and $N_{\text{Tracks}}^{\text{Event}}$.




(c) The EDP for DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5 \text{ GeV}$ outside material in the PTR.


Figure 6.3: Histograms used to calculate the EDP in the PTR. Dividing (a) by (b) produces (c). The results shown here are for DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5 \text{ GeV}$ reconstructed outside material. The number of DVs in each bin of (a) are not integers because a pileup weight is applied so that the pileup distribution of the PTR matches the pileup distribution of the MTR. This is discussed in Section 6.1.4.



(a) The EDP for DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5 \text{ GeV}$ outside material in the PTR.



(b) All events in the MTR binned by $N_{\text{b-tag}}$ and $N_{\text{Tracks}}^{\text{Event}}$.



(c) The estimated number of DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5 \text{ GeV}$ outside material in the MTR.

Figure 6.4: Histograms used to calculate the estimated number of DVs in the MTR. Multiplying (a) by (b) produces (c). The results shown here are for DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5 \text{ GeV}$ reconstructed outside material. The total number of estimated DVs with $N_{\text{Tracks}}^{\text{DV}} \geq 4$ and $m_{\text{DV}} > 5 \text{ GeV}$ in the MTR is obtained by integrating (c).


being estimated. The number of DVs observed in the PTR is shown in Figure 6.5. Due to the low statistics in the PTR, the same extended statistics method described in Section 6.1.1 is utilized here. For the extended signal region using DVs with $m_{\text{DV}} > 5 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 4$, the factor f to weight the extended statistics estimate to the signal region estimate can be obtained via Equation 6.5 as

$$f_{\text{SR Extended}} = \frac{2.03}{62.3} \approx 0.0325.$$


Initial studies comparing the performance of the track density estimation method to the performance of the jet matching method found that the track density method more accurately estimates the number of DVs both inside and outside material. It was observed that when using the jet matching method, the number of DVs that went unmatched to track jets increased with the mass of the DV, suggesting that more massive DVs are less likely to be associated to track jets and more likely to be caused by the overall track density of the event. In general, it was observed that the track density method estimated regions of higher DV mass more accurately than the jet matching method. However, the track density method also tended to underestimate regions with low DV mass relative to the jet matching method. These initial studies were conducted prior to the discovery of the significant non-collision background contamination in the MTR and the optimization of the track density method parameterization scheme. The subsequent implementation of the non-collision background veto and the optimization of the EDP parameterization greatly improved the accuracy of the track density method. Validation studies for the track density method are shown in Section 6.1.5. Considering the significant changes to the event selections and the EDP parameterization, the validity of these initial studies is questionable, so they will not be discussed further. However, the observation that the track density method more accurately estimates the number of high mass DVs and that the jet matching method more accurately estimates the number of low mass DVs motivates the development of the second alternative background method, which is ultimately used as a cross-check of the track density method.

6.1.3 Hybrid Method

The hybrid method is an attempt to capitalize on the strengths of both the jet matching method and track density method by considering the correlation between DVs and the track density of the event in addition to the correlation between DVs and



(a) Fine-binned regions



(b) Coarse-binned regions

Figure 6.5: The number of DVs in each region of the PTR with the full event and DV selections (except the mass and track multiplicity requirements) applied. (a) shows the observed DVs in the fine-binned regions and (b) shows the observed DVs in the coarse-binned regions. The bin contents are not integers because pileup weighting is applied (see Section 6.1.4). The majority of weights are of order 1.

jets. A common source of low mass DVs in the SM is the decays of b-mesons, which have non-negligible lifetimes. Though DVs from b-meson decays are not a significant source of background due to their low invariant mass, DVs from b-meson decays can be accidentally crossed by unrelated tracks or merged with nearby DVs to form DVs that pass the 1 DV SR selections. To appropriately factor in this relationship with b-meson decays, EMTopo jets are used in the hybrid method instead of track jets because the b-tagging information from EMTopo jets is readily available.

The three conceptual steps of the hybrid estimate are as follows:

1. Jet-DV matching

- An attempt is made to match DVs to EMTopo jets. The DV is matched to the highest p_T jet with $\Delta R_{\text{DV,Jet}} < 0.4$. A DV is considered unmatched if no jet satisfies $\Delta R_{\text{DV,Jet}} < 0.4$.
- All DVs that were successfully matched to jets are recorded alongside their matched jet. These DVs and jets are then removed from the event. The matched DVs and jets will be used to calculate a JDP as was done in the jet matching method.
- Any remaining unmatched DVs are attributed to the track density of the event. These DVs are then recorded alongside the track density properties of the event.

2. Calculation of JDP and EDP

- The collection of matched DVs in the PTR is used to calculate a JDP in the same way as the jet matching method. The JDP can be parameterized as a function of jet variables, such as the jet p_T or if the jet is b-tagged.
- The collection of unmatched DVs in the PTR is used to calculate an EDP in the same way as the track density method. The EDP can be parameterized as a function of event variables which describe the track density of the event, such as $N_{\text{Tracks}}^{\text{Event}}$ and $N_{\text{Track Jets}}^{\text{Event}}$.

3. Estimation


- Every jet in the MTR is assigned a weight $P(\text{DV} | \text{jet}_i)$ based on the values of the jet variables used to parameterize the JDP.

- Every event in the MTR is assigned a weight $P(\text{DV} | \text{event}_i)$ based on the values of the track density variables used to parameterize the EDP.
- The sum of weights for all jets in the MTR gives the jet matched component of the estimate, and the sum of weights for all events in the MTR gives the track density component of the estimate. The sum of the components yields the total estimate.


The default parameterization of the hybrid method uses the jet p_T and the jet b-tag for the JDP and $N_{\text{Tracks}}^{\text{Event}}$ and $N_{\text{Track Jets}}^{\text{Event}}$ for the EDP.

The number of DVs predicted outside material in the MTR by both the hybrid and track density methods is shown in Figure 6.6. Both methods estimate nearly identical numbers of DVs in the regions where $m_{\text{DV}} > 5 \text{ GeV}$. Given the comparable performance, either estimate could be chosen as the primary estimate. Though the hybrid method may appear more robust because it factors in the DV dependence on both jets and track density, the track density estimate is chosen as the primary estimate for two main reasons. The first is that there are no unmatched DVs in the PTR with $m_{\text{DV}} > 10 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 5$ when estimating with the hybrid method. This results in an EDP of 0 and the track density component not contributing to the overall estimate in the SR. Splitting the DVs into matched and unmatched categories leads to insufficient statistics in the PTR to estimate the EDP. When performing the hybrid estimate, the track density component is observed to be smaller than the jet-matched component in all CRs and VRs. The low statistics does not allow for the hybrid method to account for the production of unmatched DVs in several VRs and the SR. However, the more general track density method inclusively estimates DVs that would go both matched and unmatched in the hybrid scheme.

The second reason to use the track density method is to simplify and reduce the uncertainties. Keeping all the DVs in one histogram rather than dividing them between the matched and unmatched components can have a significant effect on the statistical uncertainty predicted by the toy method described in Section 6.1.4. Varying empty bins of a histogram within their Poisson uncertainty results in extremely large relative statistical uncertainties when integrating over the resulting histograms. Considering these two reasons, the track density method is used as the primary method for the 1 DV SR.



(a) The number of DVs predicted by the hybrid estimate for the MTR. The track density component of the estimate is parameterized by $N_{\text{Tracks}}^{\text{Event}}$ and $N_{\text{Track Jets}}^{\text{Event}}$ and the jet matched component is parameterized by the jet p_T and the jet b-tag. The uncertainties shown are the non-linearity and pileup uncertainties described in Section 6.1.4 in addition to a statistical uncertainty obtained by assigning an uncertainty of \sqrt{N} to each bin of the histograms used to calculate the estimate.




(b) The number of DVs predicted by the track density estimate for the MTR. The EDP is parameterized with $N_{\text{Tracks}}^{\text{Event}}$ and $N_{\text{b-tag}}$. The full uncertainties discussed in Section 6.1.4 are shown.

Figure 6.6: The estimated number of DVs in the MTR outside material by the (a) hybrid and (b) track density methods.


6.1.4 Uncertainties

Several sources of systematic uncertainty are assessed for this background. The first is the statistical uncertainty on the number of DVs in each CR, which is propagated to the final estimate. The uncertainties for the all-event histograms (H_{PTR} , H_{MTR}) are assessed by assigning an uncertainty of \sqrt{N} to each bin, where N is the bin content. This is sufficient for the all-event histograms where the bins are well populated and a normal distribution can be reasonably assumed. However, the DV histograms have low statistics and should be assigned Poisson uncertainties. To address this, a toy method is used which performs 10000 pseudo-experiments for the extended statistics variant of the method. This is not performed for the standard statistics version of the method due to the extremely low statistics in the regions. To estimate the uncertainty for a given VR/SR, the number of DVs in the corresponding CR is varied bin-by-bin within the statistical uncertainty. This is done by pulling a random integer from the Poisson distribution with the observed number of DVs as the mean. For example, a bin with three observed DVs would be randomly assigned a new value from the Poisson distribution with a mean of three. The resulting DV histogram is then used to calculate a new EDP. For the extended method, the calculation of the f -factor (Equation 6.4) relies on the observed number of DVs in the CR. The histogram of observed DVs for a given region often has very few events and a majority of the bins can be empty. To vary the f -factor, the number of observed DVs is not varied bin-by-bin as it is for the extended region (which has sufficient statistics for this to be a reasonable approach), but is instead pulled from the Poisson distribution with the total observed number of DVs as the mean. The estimate is then calculated with the new probability and f -factor. The standard deviation of the background estimate from all toy experiments is taken as an uncertainty on the final background estimate. The mean and standard deviation of the toy experiments for each VR and the 1 DV SR are shown in Figure 6.7.

A second systematic uncertainty is derived to factor in any residual dependence on relevant event-level quantities that were not used to parameterize the event-DV probability. To quantify this, the fraction of events in the PTR with a SR-like DV ($m_{\text{DV}} > 5 \text{ GeV}$, $N_{\text{Tracks}}^{\text{DV}} \geq 4$) is plotted as a function of three variables - $p_{\text{T}}^{2, \text{Event}}$, $N_{\text{Track Jets}}^{\text{Event}}$, and $N_{\text{EMTopo Jets}}^{\text{Event}}$. These distributions are shown in Figure 6.8. Each of these distributions is then fit to a linear function. Normalized distributions of each variable are also made for both the MTR and PTR (as in Figure 5.8). Each bin of the histograms in the PTR



(a) Fine



(b) Coarse

Figure 6.7: The mean and standard deviation of the statistically varied toy experiments for each VR and the 1 DV SR. VR1 ($2 \text{ GeV} < m_{\text{DV}} < 5 \text{ GeV}$, $N_{\text{Tracks}}^{\text{DV}} = 4$) is omitted to allow for better color scaling.

and MTR is then scaled by the value of the fitted function at the center of the bin. The relative difference between the integrals of the resulting histograms (referred to as N_{PTR} and N_{MTR}) is used as an uncertainty on the background estimate in the 1 DV SR. A flat fit of the distribution for a given variable would indicate a small dependence of the EDP on the variable, resulting in a smaller uncertainty. Conversely, if the fit shows a significant dependence of the EDP on the variable, a greater uncertainty will be assigned.

To factor in the uncertainties on the fit parameters, this procedure is repeated by varying the anti-correlated fit parameters within 1σ (shown as the red and blue lines in Figure 6.8). The largest upward and downward uncertainties from the three fits are assigned as an asymmetric uncertainty to the background estimate. The only variable with a significant contribution to this uncertainty is $N_{\text{Track Jets}}^{\text{Event}}$, resulting in an uncertainty of $^{+16\%}_{-0\%}$ on the background estimate for the 1 DV SR. None of the variables are found to contribute a downward uncertainty.

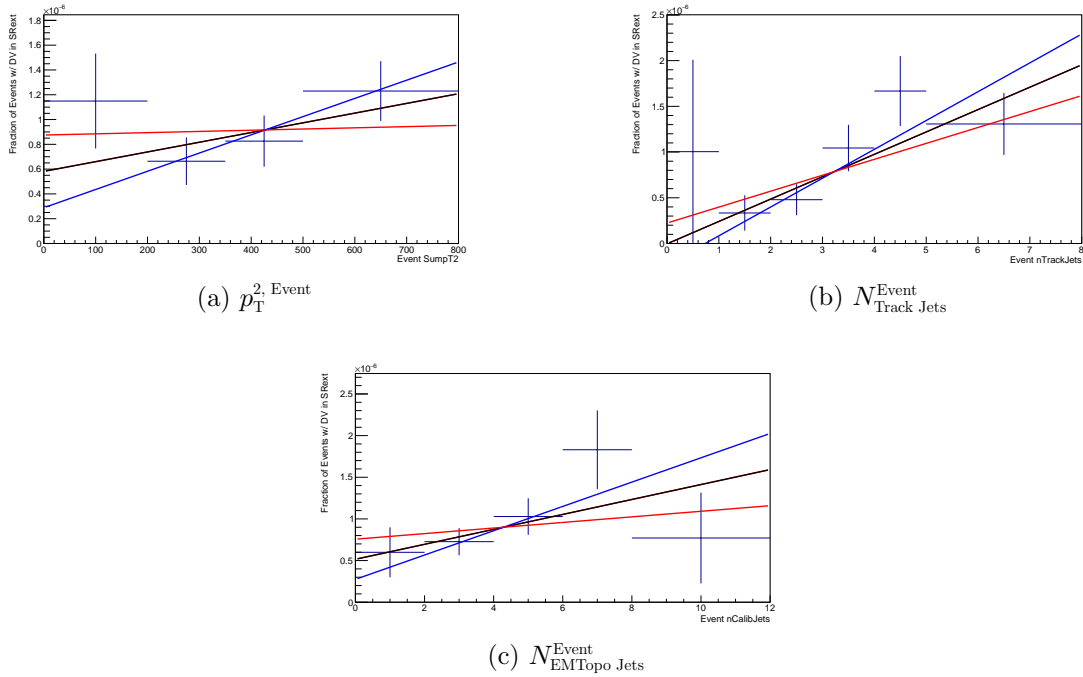


Figure 6.8: The fraction of events in the PTR with a SR-like DV ($m_{\text{DV}} > 5 \text{ GeV}$, $N_{\text{Tracks}}^{\text{DV}} \geq 4$) as a function of (a) $p_{\text{T}}^{2, \text{Event}}$, (b) $N_{\text{Track Jets}}^{\text{Event}}$, and (c) $N_{\text{EMTopo Jets}}^{\text{Event}}$. The distribution is fit to a linear function (black line), and the fit parameters are varied within 1σ (red and blue lines) to extract an asymmetric uncertainty on the estimated number of DVs.

A final uncertainty is assigned to account for differences in the pileup distributions of events in the PTR and the MTR. A plot showing the pileup distributions in the PTR and MTR as well as the ratio of the pileup in the MTR to the PTR is shown in Figure 6.9. Events in the MTR tend to have higher pileup than events in the PTR. This is expected, as E_T^{miss} triggers and selections are sensitive to increased pileup. To quantify the uncertainty from this difference, events in the PTR are weighted such that the pileup distribution matches the pileup distribution in the MTR. After applying this weight, the EDP is recalculated and used to calculate a new value for the estimate. Applying the pileup weighting procedure yields an estimate of 0.56 DVs, which is a 5% reduction in the estimated number of background DVs in the 1 DV SR. The estimated value accounting for the pileup reweighting is used as the background estimate, and the difference between the weighted and unweighted estimates is applied as an additional systematic uncertainty.




Figure 6.9: The normalized pileup distribution in the PTR (red) and MTR (blue). A ratio of the MTR to the PTR is shown on the bottom of the plot.

The total systematic uncertainty is estimated by adding these uncertainties in quadrature under the assumption that they are uncorrelated. The final inclusive background estimate for the 1 DV SR along with its uncertainties is shown in Table 6.1.

An uncertainty to account for any differences between the values estimated by the track density method, the hybrid method, and the combined background estimate method (detailed in Section 6.2) has been considered. However, no significant differences are observed when comparing the track density estimate to the hybrid

Method	Estimate	Estimate Uncertainties			Total Uncertainty
		Statistical	Non-linearity	Pileup	
Track Density	$0.56^{+0.40}_{-0.39}$	$\pm 70\%$	16% 0%	5%	+72% -70%
Hybrid	$0.60^{+0.44}_{-0.43}$	$\pm 71\%$	16% 0%	5%	+73% -71%

Table 6.1: A summary of the uncertainties on the inclusive background estimate for the 1 DV SR.

estimate or the combined estimate. One could also assess a non-closure uncertainty to account for any discrepancies between the predicted and observed number of DVs in the validation regions. VR1 is the only region where a significant difference between the predicted and observed number of DVs occurs. The background DV composition in VR1 is significantly different from that of the 1 DV SR due to being heavily dominated by b-meson decays and hadronic interactions that pass the material veto. These processes predominantly create DVs with invariant mass under 5 GeV. DVs from b-meson decays and hadronic interactions are not likely to contribute to the background in the 1 DV SR due to the high invariant mass cut of 10 GeV. Thus, no additional systematic uncertainty is assessed.

One assumption of the inclusive method when using the extended regions is that the EDP distribution for the extended signal region ($m_{\text{DV}} > 5 \text{ GeV}$, $N_{\text{Tracks}}^{\text{DV}} \geq 4$) matches that for the true signal region ($m_{\text{DV}} > 10 \text{ GeV}$, $N_{\text{Tracks}}^{\text{DV}} \geq 5$). The validity of this assumption is difficult to check due to lack of statistics in the control region. One way to study the validity of this assumption is to instead define a hyper-extended control region composed of DVs in the PTR with $m_{\text{DV}} > 5 \text{ GeV}$ and $N_{\text{Tracks}}^{\text{DV}} \geq 3$ to compare the shape of the EDP distribution to that of the extended control region. The number of DVs found in the extended control region and in the hyper-extended control region can be seen in Figure 6.10 (a) and (b), respectively. The distribution of DVs with respect to the EDP parameterization variables differs significantly between the two regions, with the hyper-extended region having the majority of its DVs in events with a high number of tracks. This is unsurprising, as low track DVs are more likely to originate from pileup. This would suggest that the assumption that the probability is similar between the signal region and the extended signal region is not a good assumption. However, if one performs the estimate using the hyper-extended control region rather than the extended region, it is seen that the central values of the estimates are consistent with each other within the statistical uncertainty. These

results are shown in Table 6.2. Because the results of the hyper-extended estimate are consistent with the extended region estimate, no additional uncertainty is assessed. Furthermore, the hyper-extended region estimate does not replace the extended region estimate because the background composition in the hyper-extended region is more dissimilar to the signal region than the extended region.

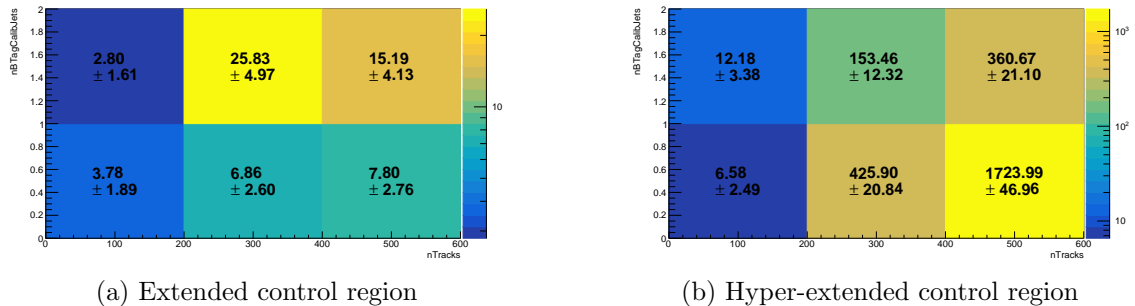


Figure 6.10: The number of DVs in the PTR in the (a) extended control region and (b) hyper-extended control region. The values shown are not integers because the DVs are weighted such that the pileup distribution in the PTR matches the distribution in the MTR.

Region	Extended Estimate	Hyper-Extended Estimate	Observed
VRLT	4.3 ± 1.2	3.2	3
VRLM	5.3 ± 1.5	3.9	4
VR2	7.0 ± 1.8	5.2	9
SR	0.56 ± 0.4	0.41	-

Table 6.2: The estimated values from the extended estimate and the hyper-extended estimate alongside the observed number of DVs in several VRs and the SR (blinded). The selections on m_{DV} and N_{Tracks}^{DV} for each region are defined in Table 5.2. The central values of the two estimate methods agree within the statistical uncertainty.

6.1.5 Validation

Two sets of validation regions are used to validate the background estimation technique:

1. Regions in data where the material map veto (described in Section 5.5.1) is inverted. These regions are enriched in DVs from hadronic interactions, but the

greater overall rate of DV production also leads to an increased number of DVs that are accidentally crossed or merged.


2. Sideband regions where the $N_{\text{Tracks}}^{\text{DV}}$ and/or the m_{DV} signal region selection is inverted. These regions are expected to have a similar background composition to the signal region while containing an acceptable amount of signal contamination.

Inside Material Validation


The number of DVs predicted for the inside material VRs is obtained by selecting DVs which pass all the baseline DV criteria outlined in Table 5.4, with the exception of the material map veto, which is instead inverted. The resulting collection of DVs is then used to calculate a new EDP for DVs produced inside material. The estimate then proceeds in the same way as the nominal estimate. Figures 6.11 (a) and (b) show respectively the estimated and observed number of DVs inside material for events passing the MTR selections. The low mass regions (VR1, VR2, VRLM) all exhibit a slight systematic overestimation of the observed number of DVs. These low mass regions are primarily populated by DVs from hadronic interactions. The dedicated hadronic interactions background estimate (outlined in Section 6.2.3) demonstrates that the contribution of these DVs to the background in the SR is expected to be negligible. Considering this, an uncertainty to cover the non-closure between the estimated and observed number of DVs in the low mass regions is not assessed. The agreement between the estimated and observed number of DVs inside material is good for DVs with $m_{\text{DV}} > 10 \text{ GeV}$.

Sideband Region Validation

The sideband regions used to validate the background estimation methods contain DVs that satisfy all the selections for the 1 DV SR, with the exception that the reconstructed DV must have $N_{\text{Tracks}}^{\text{DV}} < 5$ or $m_{\text{DV}} < 10 \text{ GeV}$. Due to the low statistics in the PTR, the EDPs for the VRs are calculated using the extended statistics method described in Section 6.1.1. The resulting estimates are then scaled by the appropriate f factor to obtain the final estimate for each region. The number of DVs estimated in the MTR by the hybrid and track density methods are shown in Figures 6.6 (a) and (b), respectively. Figures 6.12 (a) and (b) show the observed number of DVs in the MTR in the coarse-binned and fine-binned regions, respectively. Figures 6.13 and 6.14 compare the observed number of DVs to the number estimated by the track



(a) The number of DVs inside material in the MTR estimated by the track density method. The EDP is parameterized by N_{Tracks}^{Event} and N_{b-tag} .



(b) The number of observed DVs inside material in the MTR.

Figure 6.11: The estimated (a) and observed (b) number of DVs inside material in the MTR, binned by N_{Tracks}^{DV} and m_{DV} . The estimate was performed with the track density method, with the EDP parameterized by N_{Tracks}^{Event} and N_{b-tag} .


density method in the sideband regions of the 1 DV SR. The only discrepancy between the predicted and observed values occurs in VR1. As discussed in Section 6.1.4, the background composition of VR1 is not representative of the expected background in the 1 DV SR, so no additional systematic uncertainty is applied to cover this non-closure. Table 6.3 summarizes the results of the inclusive estimate validation in the sideband regions. The total number of events in the 1 DV SR estimated by the track density method is 0.56 ± 0.40 .

Name	Region		Estimate		Observed
	$N_{\text{Tracks}}^{\text{DV}}$	m_{DV}	Track Density	Hybrid	
VR1	4	[2, 5] GeV	1228 ± 203	1519 ± 252	1971
VR2	4	[5, 10] GeV	7.0 ± 1.8	7.2 ± 1.9	9
VR3	4	[10, 15] GeV	0.52 ± 0.39	0.50 ± 0.38	0
VR4	4	[15, 20] GeV	1.1 ± 0.6	1.2 ± 0.6	2
VR5	5	[5, 10] GeV	2.2 ± 0.9	2.3 ± 0.9	3
VR6	6	[5, 10] GeV	1.6 ± 0.7	1.7 ± 0.8	1
VR7	≥ 7	[5, 10] GeV	1.6 ± 0.7	1.6 ± 0.8	0
VR8	4	> 20 GeV	2.7 ± 1.0	2.9 ± 1	1
VRLM	≥ 5	[5, 10] GeV	5.3 ± 1.5	5.6 ± 1.6	4
VRLT	4	> 10 GeV	4.3 ± 1.2	4.4 ± 1.3	3
SR	≥ 5	> 10 GeV	0.56 ± 0.40	0.60 ± 0.43	Blinded


Table 6.3: Estimated number of background DVs from the track density and hybrid methods compared to the observed number of DVs in events passing the MTR. Quoted uncertainties for the track density estimate are the full uncertainties discussed in Section 6.1.4. For the hybrid estimate, the uncertainties shown are the non-linearity and pileup uncertainties described in Section 6.1.4 in addition to a statistical uncertainty obtained by assigning an uncertainty of \sqrt{N} to each bin of the histograms used to calculate the estimate.

6.2 Alternative Background Estimation

As an alternative to estimating the total background inclusively with the track density method, the expected contributions from each source of background can be estimated separately and combined to produce an estimate of the total expected background in the SR. The following sections will describe the methods used to estimate each source of background individually.



(a) Coarse-binned regions.



(b) Fine-binned regions

Figure 6.12: The observed number of DVs in the MTR, binned by $N_{\text{Tracks}}^{\text{DV}}$ and m_{DV} . The number of DVs in the SR is not shown.



(a) Low m_{DV} validation regions.(b) Low N_{Tracks}^{DV} validation regions.

Figure 6.13: Summary of the (a) low m_{DV} and (b) low N_{Tracks}^{DV} validation regions of the track density estimate. The points in black show the observed number of DVs and the red line shows the value estimated by the track density estimate. The shaded regions show the range of the full uncertainties of the track density estimate.




Figure 6.14: Summary of the merged validation regions of the track density estimate. The points in black show the observed number of DVs and the red line shows the value estimated by the track density estimate. The shaded regions show the range of the full uncertainties of the track density estimate.

6.2.1 Merged Vertices

As discussed in Section 4.2.1, the VSI secondary vertexing algorithm may merge a pair of nearby vertices if the distance significance for the pair of vertices (S , defined in Equation 4.1) is less than 10σ . It is possible that a pair of unrelated low mass vertices originating from the decays of SM LLPs or hadronic interactions could therefore be merged. Merged vertices are a source of potential background vertices because the increased mass and track multiplicity of the merged vertex could result in the vertex passing the SR selections.

The contribution to the SR background from merged vertices can be estimated with the following procedure:

1. **Derive the vertex merging rate** by comparing the distance significance distribution for pairs of DVs in the same event to the distribution for pairs of DVs in different events (where no merging could have occurred).
2. **Construct a merged vertex mass template** by artificially merging pairs of DVs from different events which satisfy the merging requirement $S < 10 \sigma$ and computing the mass of the resulting vertex.
3. **Normalize the mass template** by the vertex merging rate.

4. **Apply the complete track cleaning** to the artificially merged vertices.
5. **Integrate the mass templates** with $N_{\text{Tracks}}^{\text{DV}} \geq 5$ in the range $m_{\text{DV}} > 10 \text{ GeV}$.

Each of these steps is performed separately for each DV-track multiplicity.


Merging Rate Calculation

The first step in the estimate is the calculation of the merging rate. This is done by calculating the distance significance between pairs of DVs passing the baseline DV selections in the same event. This step is done separately for pairs of DVs with different track multiplicities (i.e. the distance significance between a pair of DVs with $N_{\text{Tracks}}^{\text{DV}} = 2$ and $N_{\text{Tracks}}^{\text{DV}} = 3$ is recorded in a separate histogram from the distance significance between a pair of DVs each with $N_{\text{Tracks}}^{\text{DV}} = 3$). If an event contains N possible pairs of DVs, then S is calculated for each pair of DVs in the event and is added to the distribution with weight $1/N$. The same procedure is then repeated using DVs from neighboring events. DV pairs from different events are referred to as ‘mixed-event’ DV pairs, and the distance significance of mixed-event DV pairs is called the ‘mixed-event distance significance’. The distance significance is calculated for a DV with respect to DVs in the following n events containing a DV passing the baseline selections, where the value of n is set to 10 to ensure adequate statistics in the distribution. The distance significance is then added to the mixed-event distance significance distribution.

For same-event DV pairs, merging cannot occur if $S > 10 \sigma$ for the pair. The mixed-event distance significance distribution is thus normalized such that the integral in the $S > 10 \sigma$ region matches that of the same-event distance significance distribution. The same-event and normalized mixed-event distribution of the distance significance squared (S^2) is shown in Figure 6.15 for pairs of DVs with two and three tracks. The shape of the mixed-event distribution matches the same-event distribution well in the $S^2 > 100$ region following the normalization. It is expected that some amount of merging has occurred between same-event pairs of DVs in the region where merging is allowed. The deficit shown of the same-event S^2 distribution relative to the mixed-event distribution in the $S^2 < 100$ region, shown in Figure 6.15 (b), confirms this expectation. The merging rate is given by

$$R_{\text{Merging}} = 1 - \frac{\int_0^{10} H_S^{\text{Same}}}{\int_0^{10} H_S^{\text{Mixed}}}, \quad (6.8)$$

where H_S^{Same} and H_S^{Mixed} are the distance significance distributions for same- and mixed-event DV pairs, respectively. Because the merging step in the VSI algorithm occurs before the track cleanings described in Section 5.5.2, the track cleanings are not applied to DVs when calculating the merging rate.



(a) S^2 for pairs of DVs with two and three tracks

(b) S^2 for pairs of DVs with two and three tracks in the $S^2 < 100$ region

Figure 6.15: The distance significance squared (S^2) for pairs of DVs in the same event (black) and mixed events (red). DVs are taken from events passing the MTR selections. The mixed-event distribution is normalized such that the integral in the $S^2 > 100$ region is the same for both the mixed-event and same-event distributions. The complete distribution is shown in (a), and (b) shows the distribution in the $S^2 < 100$ region.

Mass Template and Estimate

To produce a merged vertex mass template, pairs of baseline DVs from different events with $S < 10 \sigma$ are artificially merged. The invariant mass of the artificially merged DV is then added to a histogram based on the vertex track multiplicity. The resulting mass template is then normalized such that the total number of merged vertices is equal to the deficit of events in the same-event distance-significance distribution in the range $S < 10 \sigma$:

$$N_{\text{Deficit}} = \int_0^{10} H_S^{\text{Mixed}} - \int_0^{10} H_S^{\text{Same}}. \quad (6.9)$$

Each artificially merged vertex in the mass template is assigned an equal weight to achieve this normalization.

In the last step, track cleaning is applied to the artificially merged DVs in the mass template and their invariant mass and track multiplicity are recalculated with the passing tracks. The resulting DVs are then added to a new set of mass templates based on their post-cleaning track multiplicity with the weight assigned to them in the previous step. The final mass templates for 4- and 5- track merged vertices are shown in Figure 6.16. The expected background in the SR due to merged vertices is obtained by integrating all mass templates with $N_{\text{Tracks}}^{\text{DV}} \geq 5$ in the $m_{\text{DV}} > 10$ GeV region. An uncertainty for the estimate is derived by propagating the statistical uncertainty on the number of DV pairs in the same- and mixed- event distributions to the final estimate. The background contribution from merged vertices in the SR is estimated to be 0.18 ± 0.03 events.

6.2.2 Accidental Crossings


A DV with mass or track multiplicity below the SR thresholds can be promoted into the SR if it is accidentally crossed by an unrelated track that is then attached to the vertex during the reconstruction. If the accidentally crossing (AX) track has high p_{T} or crosses the DV at a large angle, the invariant mass of the DV can increase significantly. Background vertices caused by random crossings are part of the *accidental crossings* background. The properties of accidentally crossed DVs differ depending on the radius at which they are reconstructed, so this estimate is performed in four separate radial regions, defined in Table 6.4. The sum of the estimated background in each radial region gives the total estimated background.

Region	Minimum R_{xy} [mm]	Maximum R_{xy} [mm]
Inside Beampipe	0	25
Inside IBL	25	38
Inside L2	38	120
Inside SCT	120	300


Table 6.4: Definition of radial regions used in the accidental crossings background estimation.

The accidental crossings background is estimated in four steps:

1. **Produce a track database** from AX track candidates in data.
2. **Calculate the rate of accidental crossings** in data.



(a) Mass template for 4-track merged vertices



(b) Mass template for 5-track merged vertices

Figure 6.16: Normalized mass templates for (a) 4- and (b) 5-track merged vertices.

3. **Construct a mass template** of accidentally crossed DVs using tracks from the AX track database.
4. **Integrate the mass templates** with $N_{\text{Tracks}}^{\text{DV}} \geq 5$ in the $m_{\text{DV}} > 10 \text{ GeV}$ region to obtain an estimate for the given radial region.

The above steps are performed for each of the radial regions listed in Table 6.4.

Track Database Production

The first step of the accidental crossings estimate is to build a track database from AX track candidates in data. Tracks in data cannot be identified as accidental crossings with certainty, but if a known SM LLP decay is accidentally crossed, it is possible to identify candidate AX tracks. The K_S^0 has a mass of 498 MeV and a mean lifetime of 90 ps. The most common of its decays is to $\pi^+\pi^-$, with a branching ratio of 70%. These properties make K_S^0 decays a plentiful source of easily identifiable 2-track DVs in data.

Construction of the track database begins by iterating over all 3-track DVs in data which pass the baseline DV selections and computing the invariant mass of each pair of tracks in the DV. The masses of these so called “3-choose-2-track DVs” are binned into a histogram, which will be used when calculating the rate of accidental crossings. If any of the track pairs has a mass within 50 MeV of $m_{K_S^0}$, those tracks are flagged as a K_S^0 and the remaining track is considered to be an AX track. In the case that more than one pair of tracks in a 3-track DV have a mass consistent with $m_{K_S^0}$, the pair with the mass closest to $m_{K_S^0}$ is chosen as the K_S^0 . The AX track is then stored in the track database alongside the z and R_{xy} values of the crossed K_S^0 . Properties of the AX track, such as $\Delta\eta$ and $\Delta\phi$ with respect to the vector pointing from the primary vertex to the DV, are recorded as well to preserve the relevant kinematics of the crossing.

Crossing Factor Calculation

The probability that a DV will be accidentally crossed by an unrelated track is referred to as the *crossing factor*. The crossing factor varies with respect to the radial displacement of the DV, so separate crossing factors are calculated for the Inside L2 region, the Inside SCT region, and the combined Inside Beampipe and Inside IBL regions. A single crossing factor is calculated for the two innermost regions due to the

low available statistics. The crossing factor for a given radial region can be determined by measuring the fraction of K_S^0 candidates¹ in data that have been reconstructed with an AX track.

The number of K_S^0 decays reconstructed without an AX track in data is determined by creating a histogram of m_{DV} for all 2-track DVs passing the baseline DV selections in the MTR, as shown in Figure 6.17 (a) for 2-track DVs reconstructed with $R_{xy} < 120$ mm. In addition to the baseline DV selections, DVs must satisfy

$$\cos\left(\Delta\alpha\left(\overrightarrow{\text{PVDV}}, \vec{p}_{K_S^0}\right)\right) > 0.999, \quad (6.10)$$



where $\vec{p}_{K_S^0}$ is the 3-momentum of the K_S^0 candidate, $\overrightarrow{\text{PVDV}}$ is the vector pointing from the primary vertex to the DV, and $\Delta\alpha(\dots)$ is the 3D angle between two 3-vectors. This selection reduces the number of background DVs reconstructed near the K_S^0 mass peak by requiring that the vector sum of the DV tracks be approximately collinear with $\overrightarrow{\text{PVDV}}$. Following this selection, the K_S^0 mass peak in the 2-track m_{DV} distribution is fit to a generalized form of the Crystal Ball function [82]. The background is fit in the region outside the K_S^0 mass peak with a fourth-order polynomial function. Subtracting the fitted background from the fit of the K_S^0 mass peak and integrating the result from $m_{K_S^0} - 40$ MeV to $m_{K_S^0} + 40$ MeV provides an estimate of the number of 2-track DVs from K_S^0 decays in data.

The number of K_S^0 decays reconstructed with an AX track is estimated by applying a similar procedure to the 3-choose-2-track DV mass distribution, shown in Figure 6.17 (b) for 3-choose-2-track DVs reconstructed with $R_{xy} < 120$ mm. These DVs must satisfy the same requirements as the 2-track DVs from the previous step. A fourth-order polynomial is again used to fit the background, and the K_S^0 mass peak is fit with a Gaussian distribution. The number of accidentally crossed K_S^0 candidates is obtained by subtracting the background fit from the K_S^0 peak fit and integrating from $m_{K_S^0} - 40$ MeV to $m_{K_S^0} + 40$ MeV.

The crossing factor for each radial region is then calculated as

$$f_C = \frac{N_{3\text{-track}}^{K_S^0}}{N_{2\text{-track}}^{K_S^0} + N_{3\text{-track}}^{K_S^0}}, \quad (6.11)$$

¹Due to the probabilistic nature of fundamental particles and their decays, the true origin of these DVs cannot be known with 100% certainty. As such, DVs consistent with K_S^0 decays are referred to as K_S^0 candidates.

(a) 2-track m_{DV} distribution.(b) 3-choose-2-track m_{DV} distribution.Figure 6.17: Mass distributions for (a) 2- and (b) 3-choose-2-track DVs for DVs with $R_{xy} < 120 \text{ mm}$.

where $N_{3\text{-track}}^{K_S^0}$ and $N_{2\text{-track}}^{K_S^0}$ are the number of K_S^0 candidates in data with and without an AX track, respectively. Table 6.5 shows the measured crossing factors for each of the radial regions used in the estimate. The uncertainties on the fits used to calculate $N_{2\text{-track}}^{K_S^0}$ and $N_{3\text{-track}}^{K_S^0}$ are propagated to assess an uncertainty on the crossing factor.

Radial Region	Crossing Factor (f_C)
Inside BP Inside IBL	$2.00 \times 10^{-4} \pm 5.04 \times 10^{-5}$
Inside L2	$8.83 \times 10^{-5} \pm 1.35 \times 10^{-5}$
Inside SCT	$2.04 \times 10^{-4} \pm 2.92 \times 10^{-5}$

Table 6.5: K_S^0 crossing factors for each radial region. A single crossing factor is calculated for the two inner regions due to the lack of statistics in those regions.

Mass Template Construction and Estimate

For a given track multiplicity n , an $(n + 1)$ -track mass template for accidentally crossed DVs is constructed by artificially attaching AX tracks from the track database to n -track DVs in the MTR that pass the baseline DV selections. Separate mass templates are built for each radial region of the estimate. For each n -track DV in the MTR which passes the baseline DV selections, a random track from the same radial region is selected from the track database to be artificially attached to the DV. The chosen track is initially required to originate from the same z region as the target DV, but this requirement is relaxed if no tracks in the database share both a radial and z bin with the DV. The mass of the $(n + 1)$ -track DV is then recalculated and added to the $(n + 1)$ -track mass template of the corresponding radial region. After attaching an AX track to all baseline DVs in the MTR, the mass templates for each radial region are scaled by the corresponding crossing factor. The contribution of accidentally crossed DVs in the signal region is then estimated by integrating all $(n + 1)$ -track mass templates in the $m_{\text{DV}} > 10 \text{ GeV}$ region for all $n \geq 4$. The uncertainties on the crossing factors are combined with the statistical uncertainties of the unnormalized mass templates to produce an uncertainty on the final estimate. The background contribution from accidentally crossed vertices in the SR is estimated to be 0.76 ± 0.15 events.

6.2.3 Hadronic Interactions

The hadronic interactions component of the alternative background estimate is used to predict the number of DVs in the SR that come from SM processes that are unaffected by reconstruction effects. The majority of such ‘true DVs’ originate from strongly interacting SM particles interacting with the dense material of the ID. These are referred to as *hadronic interaction (HI) DVs*. True DVs in the ID can also be produced by the decays of SM LLPs, such as b-hadrons. Both processes tend to produce low mass DVs composed of highly collimated tracks. Enforcing the material map veto described in Section 5.5.1 and the PV-DV separation cut described in Section 5.5 removes the overwhelming majority of vertices belonging to this category of background.

Certain components of the detector, such as cables and gas lines, are not included in the GEANT4 model of the detector geometry, so DVs in some of these less dense regions may not be vetoed by the material map veto. Additionally, some HI DVs may pass the material map veto due to their positions being poorly reconstructed. Truth studies of simulated background samples performed by the DV + Jets analysis found that these residual HI DVs are the dominant background component in the low mass region of the m_{DV} distribution [2]. For a given track multiplicity, the mass distribution for HI DVs in simulation exhibits a sharp linear increase at low mass before peaking near 2 GeV. The high mass region of the distribution then falls off exponentially from the low mass peak. This distribution can be modeled with the function

$$N_{\text{HI}}(m_{\text{DV}}) = \frac{1}{\frac{1}{C(m_{\text{DV}}-b)} + e^{\frac{m_{\text{DV}}-B}{l}}}, \quad (6.12)$$

where the parameters C and b parameterize the linear rise in the low mass region, and B and l are used to parameterize the exponential decay of the high mass tail. Mass templates for each track multiplicity are created by fitting the mass distribution of baseline DVs in data in the region with $m_{\text{DV}} < 10$ GeV. The contribution of HI DVs to the SR background is obtained by integrating the fit function in the $m_{\text{DV}} > 10$ GeV region for all mass templates with $N_{\text{Tracks}}^{\text{DV}} \geq 5$. For each track multiplicity, an uncertainty on the estimated number of DVs is calculated by varying each of the fit parameters within one standard deviation and recalculating the estimate. This procedure is repeated many times, and the standard deviation of the resulting estimates is taken as an uncertainty on the nominal estimate. The background


contribution from hadronic interaction vertices in the SR is estimated to be less than 10^{-4} events, and is therefore negligible in comparison to the accidental crossing and merged vertex components.

6.2.4 Combined Estimate


The total background is estimated by adding the individual background estimates for the accidental crossings, merged vertices, and hadronic interactions components. The uncertainties for each of the components are assumed to be uncorrelated, so they are added in quadrature to produce the uncertainty for the final estimate. This procedure produces a final estimated background in the SR of 0.94 ± 0.15 . Figures 6.18 (a) and (b) show the estimated background from the combined method in the fine-binned and coarse-binned regions, respectively. Table 6.6 shows a comparison of the background estimated by the track density and combined methods in the validation regions and the 1 DV SR. The observed number of DVs for each unblinded region is also shown. Table 6.7 summarizes the predicted background in the 1 DV SR from each of the background estimation methods. The track density, hybrid, and combined background estimates all agree within statistical uncertainty.

Name	Region		Estimate		Observed
	$N_{\text{Tracks}}^{\text{DV}}$	m_{DV}	Track Density	Combined	
VR1	4	[2, 5] GeV	1228 ± 203	2000 ± 360	1971
VR2	4	[5, 10] GeV	7.0 ± 1.8	8.8 ± 2.0	9
VR3	4	[10, 15] GeV	0.52 ± 0.39	0.64 ± 0.09	0
VR4	4	[15, 20] GeV	1.1 ± 0.6	0.64 ± 0.09	2
VR5	5	[5, 10] GeV	2.2 ± 0.9	4.6 ± 1.3	3
VR6	6	[5, 10] GeV	1.6 ± 0.7	1.5 ± 0.9	1
VR7	≥ 7	[5, 10] GeV	1.6 ± 0.7	-	0
VR8	4	> 20 GeV	2.7 ± 1.0	1.6 ± 0.2	1
VRLM	≥ 5	[5, 10] GeV	5.3 ± 1.5	6.1 ± 1.6	4
VRLT	4	> 10 GeV	4.3 ± 1.2	2.9 ± 0.2	5
SR	≥ 5	> 10 GeV	0.56 ± 0.40	0.94 ± 0.15	-

Table 6.6: Estimated number of background DVs from the track density estimate and the combined background estimate in the validation regions and signal region for the 1 DV SR.



(a) Fine-binned



(b) Coarse-binned

Figure 6.18: The estimated background from merged vertices, accidental crossings, and hadronic interactions in signal and validation regions with (a) fine binning and (b) coarse binning. Both plots are binned in m_{DV} and N_{Tracks}^{DV} .

Method	Estimate
Inclusive - Track Density	$0.56^{+0.40}_{-0.39}$
Inclusive - Hybrid	$0.60^{+0.44}_{-0.43}$
Combined	0.94 ± 0.15
Merged Vertex	0.18 ± 0.03
Accidental Crossings	0.76 ± 0.15
Hadronic Interactions	$1.8 \times 10^{-5} \pm 5.9 \times 10^{-5}$

Table 6.7: Predicted background in the 1 DV SR from each of the background estimation methods.

6.3 Systematic Uncertainties

Estimated signal yields derived from the simulated signal samples described in Section 5.3.2 are used to interpret the results of this search. The predicted signal yields are subject to several sources of uncertainty that must be accounted for when comparing them to the observed data in the signal region. These uncertainties are roughly divided into two categories:

- *Experimental uncertainties* are derived to quantify performance differences between the simulated samples and data. These also include uncertainties on the integrated luminosity of the data set.
- *Theoretical uncertainties* quantify uncertainties on the modeling of the signal process.

The uncertainties considered for the benchmark gluino model are discussed in the following sections.

6.3.1 Tracking and Vertexing Uncertainties

Differences in the performance of LRT and secondary vertexing between simulation and data are a key source of uncertainty on the estimated signal yields. The first step in assessing an uncertainty due to these differences is to estimate the efficiency of LRT in simulation relative to data. This is done by comparing the number of reconstructed K_S^0 DVs in data to simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples. Candidate 2-track K_S^0 DVs are identified with the same approach and criteria used in the accidental crossings background estimate. To properly normalize the number of candidate K_S^0 DVs in the

simulated sample, the candidate K_S^0 yield in the simulated samples is scaled such that the number of candidate K_S^0 DVs with $R_{xy} < 30$ mm matches the number observed in data. The low radius region is chosen because K_S^0 decays at low radii typically have both tracks reconstructed by the standard tracking algorithm, whose efficiency in data and simulation is well understood [83]. Figure 6.19 shows the K_S^0 candidate DV yield in 2016 data (black) and simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples with the previously discussed normalization (red) as a function of DV R_{xy} . The simulated samples have additional pileup interactions from the 2016 data-taking period overlaid on them. The bottom portion of the plot shows the ratio of the yield in data to simulation. Figures 6.20 and 6.21 show the same comparison for 2017 and 2018 data and simulation, respectively. An uncertainty on the tracking efficiency in each radial bin is assigned as

$$\sigma_\epsilon = (1 - \sqrt{\epsilon_K}) \oplus 0.017, \quad (6.13)$$

where ϵ_K is the ratio of the K_S^0 yield in data to simulation for bin K . An additional uncertainty of 1.7% is added in quadrature to account for the uncertainty on the standard tracking efficiency [83]. The uncertainty on the tracking efficiency ranges between 2% and 22% depending on the radial bin and the year during which the data was collected.

The uncertainty on the tracking efficiency is propagated to the estimated signal yield by performing a track-killing study. For a given signal sample, each DV in the sample has tracks randomly removed at a rate corresponding to the tracking uncertainty in the radial bin for the data-taking year the sample is modeling. The invariant mass and track multiplicity of the DV are then recalculated. Once this process has been applied to all DVs in the event, the signal yield is recalculated. To account for the probabilistic nature of the track removal, this process is performed multiple times, with the track-killed yield given by the average of all of the runs. The difference between the track-killed and nominal yields is assigned as a symmetric uncertainty on the signal yield. A maximum uncertainty of 10% is assessed for the gluino signal samples. This uncertainty is larger for samples with smaller mass splittings between the gluino and neutralino because the DVs in these samples tend to have lower track multiplicity. In these cases, it is more likely that removing one track from a DV will cause it to fail the signal region selections.




Figure 6.19: Distribution of the number of K_S^0 candidate DVs as a function of the radial displacement of the DV. The yield in 2016 data is shown in black and the yield in simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown in red. The yield in the simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is normalized such that the yield in the first bin ($R_{xy} < 30$ mm) matches the yield in data. A plot showing the ratio of the yield in 2016 data to the normalized yield in the $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown at the bottom of the figure.

6.3.2 Jet and E_T^{miss} Uncertainties

The efficiency of the signal region selections is affected by the modeling of E_T^{miss} in the simulated samples. In the case of the gluino samples, the largest components of the E_T^{miss} are associated with jets and the soft term. Differences between jets in data and simulation are therefore important to consider due to their impact on the E_T^{miss} reconstruction. Four sources of uncertainty are considered - the jet energy scale (JES), the jet energy resolution (JER), the scale of the E_T^{miss} soft term, and the resolution of the E_T^{miss} soft term. The effect of the uncertainties on the reconstructed E_T^{miss} is obtained by producing samples where nuisance parameters associated with each uncertainty source are varied. These variations are handled by central tools developed by the ATLAS jet/ E_T^{miss} combined performance group. The signal yields are then recalculated for each variation and their difference with respect to the nominal yields is taken as an uncertainty. The uncertainties from each variation are assumed to be uncorrelated and are added in quadrature to obtain a final uncertainty on the signal yield from the modeling of E_T^{miss} and jets in simulation. The maximum observed




Figure 6.20: Distribution of the number of K_S^0 candidate DVs as a function of the radial displacement of the DV. The yield in 2017 data is shown in black and the yield in simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown in red. The yield in the simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is normalized such that the yield in the first bin ($R_{xy} < 30$ mm) matches the yield in data. A plot showing the ratio of the yield in 2017 data to the normalized yield in the $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown at the bottom of the figure.

uncertainty from these effects is 16%, but the majority of samples in the gluino grid are assigned an uncertainty of less than 5%.

6.3.3 E_T^{miss} Trigger Uncertainties

The efficiency of an E_T^{miss} trigger with respect to the offline E_T^{miss} is defined as the fraction of events with a given offline E_T^{miss} that pass the E_T^{miss} trigger selections. To measure the efficiency of an E_T^{miss} trigger in data, a highly efficient reference trigger is needed so that the measurement of both the numerator and denominator of the efficiency is unbiased. As discussed in Section 4.3.5, muons are not included in the calculation of E_T^{miss} at the trigger level. Muon triggers in ATLAS are extremely efficient, and their behavior is well understood, making muon-triggered events an excellent reference to understand the efficiency of E_T^{miss} triggers. Events containing muons are indistinguishable from events containing true E_T^{miss} to the E_T^{miss} trigger. The offline E_T^{miss} in muon-triggered events can be made to appear as it would in events with true E_T^{miss} by treating muons as invisible when calculating the offline E_T^{miss} .




Figure 6.21: Distribution of the number of K_S^0 candidate DVs as a function of the radial displacement of the DV. The yield in 2018 data is shown in black and the yield in simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown in red. The yield in the simulated $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is normalized such that the yield in the first bin ($R_{xy} < 30$ mm) matches the yield in data. A plot showing the ratio of the yield in 2018 data to the normalized yield in the $Z \rightarrow \nu\bar{\nu} + \text{jets}$ samples is shown at the bottom of the figure.

For the purpose of this search, this version of the E_T^{miss} is referred to as METNoMu. Muon-triggered events can therefore accurately represent events containing true E_T^{miss} both at the trigger level and offline, while remaining unbiased due to being collected by highly efficient muon triggers. Figure 6.22 shows the efficiency of several E_T^{miss} triggers used in the 2018 with respect to METNoMu for $Z \rightarrow \mu^+\mu^-$ events in data [68]. The triggers reach an efficiency plateau at an offline E_T^{miss} of approximately 200 GeV. The region with $E_T^{\text{miss}} < 200$ GeV where the trigger is not fully efficient is referred to as the *trigger turn-on region*. Events in the MTR are required to have $E_T^{\text{miss}} > 150$ GeV, meaning that some fraction of events are collected with the trigger operating in its turn-on region.

It is possible that the modeling of the E_T^{miss} trigger turn-on regions in simulation is inconsistent with the observed efficiency turn-on in data. To check this, the efficiency of the trigger in the turn-on region is plotted as a function of METNoMu for events in data and simulated $Z \rightarrow \mu^+\mu^-$ events. This is shown for the lowest unprecaled E_T^{miss} trigger active in 2017 in Figure 6.23. Events in Figure 6.23 are required to pass a muon trigger and have $\text{MET}_{\text{LocHadTopo}} > 180$ GeV as required by the DRAW_RPVLL filter. The




Figure 6.22: E_T^{miss} trigger efficiencies for several E_T^{miss} triggers used in 2018 data-taking with respect to the offline E_T^{miss} in $Z \rightarrow \mu^+ \mu^-$ events in data [68]. Muons are treated as invisible when calculating the offline E_T^{miss} for these events in order to more accurately represent events containing sources of true E_T^{miss} .

ratio of data to simulation is within 2% of unity for all bins with $E_T^{\text{miss}} > 150$ GeV. The quality of the agreement between data and simulation does not change when the additional requirement that a baseline DV be reconstructed in the event is imposed. The E_T^{miss} trigger efficiency in muon-triggered events can also be plotted as a function of the offline E_T^{miss} with muons treated as visible. Figure 6.24 shows the efficiency for the lowest unrescaled E_T^{miss} trigger active in 2018 with respect to offline E_T^{miss} for events which pass a muon-trigger and have $\text{MET}_{\text{LocHadTopo}} > 180$ GeV. The points in red show the trigger efficiency for simulated gluino samples while the points in blue show the efficiency in data. The ratio of the efficiency in data and simulated gluino samples deviates from unity by less than 3%. The uncertainty associated with the modeling of the E_T^{miss} trigger turn-on efficiency in the gluino signal samples is small compared to the uncertainties on the gluino pair production cross section, so this uncertainty is not applied in the limit setting process described in Chapter 7.




Figure 6.23: Efficiency of the lowest unrescaled E_T^{miss} trigger active in 2017 with respect to METNoMu for 2017 data (blue) and simulated $Z \rightarrow \mu^+ \mu^-$ events (red). The bottom plot shows the ratio of the efficiency in data to the efficiency in simulated events. The agreement in the region used by the analysis ($E_T^{\text{miss}} > 150$ GeV) is within 2%.

6.3.4 Theoretical Uncertainties

The production cross sections for the gluino samples (described in Sections 2.5 and 5.3) and their uncertainties have been calculated at approximate next-to-next-to-leading order with next-to-next-to-leading logarithmic accuracy in accordance with the PDF4LHC15 recommendations [84–91]. The uncertainty on the gluino production cross section ranges from 8% for 400 GeV gluinos up to 48% for 3000 GeV gluinos.

Additional uncertainties on the modeling of the signal process arise from the choice of PDF and the factorization and renormalization scales. During MC event generation, uncertainties from the PDF and scale variations are stored as on-the-fly event weights. The weights corresponding to the factorization and renormalization scale uncertainties are obtained by varying the scale up and down by a factor of two. These weights are then applied to the generated events and the signal yield is assessed. The largest deviation from the nominal yield is assigned as a systematic uncertainty to cover variations of these scales. A set of weights associated with variations of the chosen PDF is also calculated for each event. The signal efficiency for each variation is




Figure 6.24: Efficiency of the lowest unprescaled E_T^{miss} trigger active in 2018 with respect to offline E_T^{miss} for 2018 data (blue) and simulated gluino events (red). The bottom plot shows the ratio of the efficiency in data to the efficiency in simulated events. The maximum deviation between data and simulation in the region used by the analysis ($E_T^{\text{miss}} > 150$ GeV) is less than 3%.

calculated, and the deviations from the nominal signal efficiency are used to assess an uncertainty, following the PDF4LHC15 recommendations [91].

6.3.5 Additional Uncertainties

When simulated events are generated, additional pileup interactions are overlaid on the event to more accurately represent the detector environment and conditions inside ATLAS. However, the simulated pileup conditions are not identical to the true pileup conditions for a given data-taking period. To correct for this, simulated signal events are weighted such that the distribution of the number of pileup interactions in the samples matches the pileup distribution of the corresponding data period in a process called *pileup reweighting*. The yields from the pileup reweighted samples are taken as the nominal yields, and the difference between the weighted and unweighted yields is taken as a systematic uncertainty. The uncertainty from pileup reweighting for the gluino signal samples is found to be negligible compared to the cross section and tracking uncertainties.

The integrated luminosity of the 2016-2018 data set is used to scale the predicted signal yields. The uncertainty on the integrated luminosity for this period is 0.84% [37]. This uncertainty is applied to the luminosity-normalized signal yields.

Chapter 7

Displaced Vertices + E_T^{miss} : Results and Interpretation

The results of this analysis are determined by comparing the observed number of events in the signal region of the data set to the estimated number of background events in the signal region. After unblinding the data set, a single event is observed in the 1 DV SR, compared to an estimated SM background of $0.56_{-0.39}^{+0.40}$ from the track density estimation method described in Section 6.1.2. This observation is also consistent with the predicted background from the two alternative background estimation methods. A comparison of the observed data and the estimated background is shown in Table 7.1.

Estimation Method	Estimate	Observed
Inclusive - Track Density	$0.56_{-0.39}^{+0.40}$	1
Inclusive - Hybrid	$0.60_{-0.43}^{+0.44}$	
Combined	0.94 ± 0.15	

Table 7.1: Comparison of the estimated number of background events in the 1 DV SR from the various methods described in Chapter 6 to the observed number of events in the 1 DV SR in data.

The single event in data that passes the 1 DV SR selections was recorded in 2017. Table 7.2 summarizes the event-level variables of the signal region event. Comparing these values to the distributions in Section 5.4.3, it can be seen that each of these variables falls into the typical ranges for events in the MTR. Tables 7.3 and 7.4 list the properties of the six EMTopo jets and four track jets present in the event, respectively. Each of the four track jets in the event is reconstructed within the $\Delta R < 0.4$ cone of an EMTopo jet.

Variable	Value
Event number	1534801244
Run number	338897
$N_{\text{Track Jets}}^{\text{Event}}$	4
$N_{\text{EMTopo Jets}}^{\text{Event}}$	6
$N_{\text{b-tag}}$	2
Leading jet p_{T}	232 GeV
$E_{\text{T}}^{\text{miss}}$	211 GeV
$E_{\text{T}}^{\text{miss}} \phi$	1.54
$N_{\text{Tracks}}^{\text{Event}}$	487
# of pileup interactions	46

Table 7.2: Event-level variables of the 1 DV SR event.

Jet index	p_{T} [GeV]	η	ϕ	b-tag
1	232	2.31	-0.95	False
2	166	-1.36	1.73	True
3	153	0.69	-2.38	True
4	75.1	0.75	-1.88	False
5	29.9	-2.94	1.03	False
6	20.8	-2.29	-0.52	False

Table 7.3: Properties of the EMTopo jets in the 1 DV SR event. The jet η and ϕ are defined with respect to the primary vertex.

The properties of the DV passing the 1 DV SR selections are shown in Table 7.5, and Table 7.6 lists the properties of the DV’s constituent tracks. The tracks at index one, three, four and five of Table 7.6 are highly collimated in η and ϕ , but the track at index two has significantly different η and ϕ than the other four tracks. Given the large η and ϕ differences between track two and the other tracks, it is likely that track two is an accidentally crossing track. The η and ϕ of the DV with respect to the primary vertex place it a distance $\Delta R_{\text{DV,Jet}} = 0.13$ from the EMTopo jet at index three in Table 7.3. The fact that the DV is reconstructed at a low radius and within the jet cone of a b-tagged EMTopo jet makes this vertex consistent with the decay of a SM b-hadron. The m_{DV} distribution for DVs from b-hadron decays is strongly peaked at $m_{\text{DV}} < 5$ GeV, making it extremely unlikely for a DV from a pure b-hadron decay to pass the invariant mass cut of the 1 DV SR. However, the inclusion of an accidentally

Jet index	p_T [GeV]	η	ϕ	Corresponding EMTopo jet index
1	117	2.30	-0.96	1
2	61.8	0.68	-2.38	3
3	32.5	0.60	-1.93	4
4	14.2	-1.31	1.74	2

Table 7.4: Properties of the track jets in the 1 DV SR event. The jet η and ϕ are defined with respect to the primary vertex.

crossing track at a high angle could sufficiently increase m_{DV} to push the DV into the signal region. Indeed, if the invariant mass of the vertex is recalculated without including track two, the vertex has an invariant mass of 2.09 GeV, which is far below the requirement for the 1 DV SR. As shown in Table 6.7, the majority of the predicted background in the 1 DV SR comes from accidentally crossed DVs. The observed event appears to be consistent with this prediction. However, due to the probabilistic nature of fundamental particles and their decays, this is purely speculation which cannot be confirmed.

Variable	Value
m_{DV}	15.1 GeV
N_{Tracks}^{DV}	5
$N_{\text{Sel. Tracks}}^{DV}$	3
$N_{\text{LRT Tracks}}^{DV}$	0
R_{xy}	5.88 mm
η	2.84
ϕ	-2.42
η relative to PV	0.75
ϕ relative to PV	-2.50

Table 7.5: Properties of the DV in the 1 DV SR event.

7.1 Statistical Analysis

In the absence of a significant excess of events in the signal region, upper limits on the gluino pair production cross section are set by performing a statistical analysis of the observed data and the predicted signal yields. These upper limits are used to exclude regions of the model phase space as a function of the model parameters,

Track index	p_T [GeV]	η	ϕ	Selected / Attached
1	5.80	0.71	-2.54	Selected
2	2.09	2.34	-0.86	Selected
3	4.67	0.68	-2.55	Selected
4	4.85	0.67	-2.38	Attached
5	4.53	0.84	-2.40	Attached

Table 7.6: DV track properties for the DV found in the 1 DV SR.

namely the mass of the gluino ($m_{\tilde{g}}$), the mass of the lightest neutralino ($m_{\tilde{\chi}_1^0}$), and the mean proper lifetime of the gluino (τ).

The results of this analysis are performed with a hypothesis test. When searching for a hypothetical process, the results of the search must be compared to the expected results of the *null hypothesis*. In the case of this analysis, the null hypothesis is given by the set of known Standard Model processes, and the expected results of the null hypothesis are referred to as the background. The null hypothesis is often referred to as the *background-only hypothesis*. The presence of some additional BSM physics in addition to the SM is referred to as the *signal hypothesis*. A hypothesis can be either accepted or rejected based on a quantified level of disagreement between the hypothesis and the observed data. The variable used to quantify the disagreement between the prediction and the observed data is referred to as the *test statistic*.

The test statistic used by this analysis is derived from the likelihood function

$$L(\text{Data}|\mu, \boldsymbol{\theta}) = \prod_{i=1}^{N_{\text{bins}}} \frac{(\mu s_i + b_i)^{n_i}}{n_i!} \cdot e^{-(\mu s_i + b_i)}, \quad (7.1)$$

where μ is the signal strength, $\boldsymbol{\theta}$ is the set of nuisance parameters corresponding to the uncertainties on the predicted background and signal yields, N_{bins} is the number of bins used in the fit, and s_i and b_i are the expected number of signal and background events in bin i , respectively [92]. A signal strength of $\mu = 0$ corresponds to the background-only hypothesis and a signal strength of $\mu = 1$ corresponds to the signal hypothesis. More intuitively, μ is a scale factor on the predicted cross section of the signal process. In the case of the 1 DV SR, only a single bin, the yield in the 1 DV SR, is used when calculating the likelihood. The profile likelihood ratio is defined as

$$\lambda(\mu) = \frac{L(\mu, \hat{\boldsymbol{\theta}})}{L(\hat{\mu}, \hat{\boldsymbol{\theta}})}, \quad (7.2)$$

where $\hat{\theta}$ is the value of θ that maximizes L for the specified value of μ , and $\hat{\theta}$ and $\hat{\mu}$ are the maximum likelihood estimators of θ and μ [92]. The numerator of Equation 7.2, known as the profile likelihood, is only a function of the signal strength μ . This makes the profile likelihood ratio also only a function of the signal strength. The denominator of the profile likelihood ratio is the maximized likelihood function, which ensures that $\lambda(\mu) \leq 1$ [92].

The test statistic used to interpret the results of this analysis is defined as

$$\tilde{q}_\mu = \begin{cases} -2 \ln \frac{L(\mu, \hat{\theta}(\mu))}{L(0, \hat{\theta}(0))} & \hat{\mu} < 0 \\ -2 \ln \frac{L(\mu, \hat{\theta}(\mu))}{L(\hat{\mu}, \hat{\theta})} & 0 \leq \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu \end{cases} \quad (7.3)$$

The value of \tilde{q}_μ quantifies the level of agreement between the hypothesis and the observed data. Equation 7.3 is valid in the scenario where the presence of a signal can only increase the expected number of events in a bin. Smaller values of \tilde{q}_μ correspond to a greater degree of compatibility between the hypothesized value of μ and the observed data, while greater values of \tilde{q}_μ correspond to a higher degree of incompatibility with the hypothesized value of μ . When setting upper limits on the signal strength, an observation of $\hat{\mu} > \mu$ in the data should not be interpreted as being incompatible with the hypothesized value of μ , so the value of \tilde{q}_μ is set to zero in this scenario. In other words, only lower values of μ should be considered as alternative hypotheses when setting an upper limit on μ . A *p-value* for the signal-plus-background hypothesis ($\mu = 1$) is calculated from the test statistic through the formula

$$p_{s+b} = \int_{\tilde{q}_{\text{obs}}}^{\infty} f(\tilde{q}|s+b)d\tilde{q}, \quad (7.4)$$

where $f(\tilde{q}|s+b)$ is the probability density function for \tilde{q}_μ under the assumption of the signal-plus-background hypothesis, and $\tilde{q}_{\mu, \text{obs}}$ is the value of the test statistic observed in data. The *p-value* for the background-only hypothesis is given by

$$p_b = \int_{-\infty}^{\tilde{q}_{\text{obs}}} f(\tilde{q}|b)d\tilde{q}. \quad (7.5)$$

The probability density function of the test statistic for different values of μ is obtained using either pseudo-experiments or asymptotic formulae [92]. In the case of small expected signal and background yields, the asymptotic approximations for the distribution of \tilde{q}_μ break down. The expected number of background events in

the 1 DV SR is small, as is the predicted signal yield for many of the samples in the benchmark gluino model. Considering this, several values of μ are chosen for each sample, and one hundred thousand pseudo-experiments are run for each value of μ . In each pseudo-experiment, each of the nuisance parameters is varied within its uncertainty, and the test statistic is recalculated. The probability density function $f(\tilde{q}_\mu|\mu)$ is determined from these pseudo-experiments, and the p-value for the observed data for the given value of μ is calculated. The lower limits on the gluino mass calculated with the pseudo-experiment method differ by 10 GeV or less from the limits calculated using asymptotic approximations. Because the difference between the methods is small, the limit plots shown in Section 7.2, which were produced using the asymptotic approximation method, were not updated using the results from pseudo-experiments.

Upper limits on the benchmark model are set using the CL_s method. The CL_s value is defined as

$$\text{CL}_s = \frac{p_{s+b}}{1 - p_b}, \quad (7.6)$$

where p_{s+b} is the p-value for the signal-plus-background hypothesis, and p_b is the p-value for the background-only hypothesis [93]. A signal hypothesis at a given signal strength μ is excluded at 95% confidence level if the CL_s value is less than 0.05. When setting limits using only p_{s+b} , it is possible to exclude regions of parameter space that the analysis is not sensitive to. This can occur in scenarios where the probability density functions for the test statistic of the background-only hypothesis and the signal hypothesis are not well separated, such as when the expected number of signal events is much less than the expected number of background events and the observed number of events in data has a downward fluctuation. In such a scenario, p_{s+b} will be small and p_b will be large. Using p_{s+b} alone could result in the exclusion of the signal model. However, dividing p_{s+b} by $1 - p_b$ will increase the CL_s value and cause the signal to not be excluded. The CL_s method leads to more conservative limits than using p_{s+b} , but it avoids the issue of excluding signals that the analysis is not sensitive to.

7.2 Interpretation

Upper limits on the gluino pair production cross section are calculated with the pyhf framework [94, 95]. A single-bin fit of the yield in the 1 DV SR is used to extract the CL_s values for each sample. Expected limits are obtained using the Asimov data set, where the observed data set is assumed to exactly match the data set predicted by the background-only hypothesis [92]. Figure 7.1 shows the 95% confidence level upper limit on the gluino pair production cross section for the benchmark split-supersymmetry model as a function of the gluino mass for a fixed mass splitting of 100 GeV between the gluino and lightest neutralino. The upper limits on the cross section are shown for gluino mean proper lifetimes of 0.01 ns, 0.1 ns, 1 ns, and 10 ns. The predicted cross section as a function of gluino mass is shown by the solid black line. For each lifetime, the solid line shows the observed limit while the dashed line shows the expected limit. The colored band around the expected limit shows the $\pm 1\sigma$ variations of the expected limit. The strictest limit on the cross section for a given gluino mass is for gluinos with a mean proper lifetime of 0.1 ns. This is expected, as samples with longer lifetimes are more likely to have the gluino decay outside the fiducial volume, and samples with shorter lifetimes are likely to be rejected by the PV-DV separation cut. For a given lifetime, the observed lower limit on the gluino mass assuming the cross section predicted by the theory and a fixed mass splitting of 100 GeV is given by the gluino mass where the observed limit and the cross section intersect.

Figure 7.2 shows the expected and observed 95% confidence level exclusion limits on the gluino mass as a function of the mean proper lifetime of the gluino for a fixed neutralino mass of 100 GeV. Figures 7.3 and 7.4 show the same limits for scenarios with fixed mass splittings of 100 GeV and 30 GeV, respectively. The dashed black line and the yellow band show the expected limit and its $\pm 1\sigma$ variations. The solid red contour shows the observed limit, while the observed limits when the gluino pair production cross section is varied up and down by its theoretical uncertainty are shown by the dashed red line. Gluino masses and lifetimes that lie under the curve are excluded at 95% confidence level. The most stringent limits are set for lifetimes of approximately 0.1 ns. As the mass splitting between the gluino and neutralino gets smaller, the limits on the gluino mass for a given lifetime decrease. Scenarios with smaller mass splittings have less intrinsic E_T^{miss} than scenarios with larger mass splittings, so the efficiencies of these scenarios relative to those with larger mass splittings are negatively impacted by the offline E_T^{miss} requirement. This can be seen




Figure 7.1: Upper limits on the gluino pair production cross section for the split-supersymmetry benchmark model as a function of gluino mass, assuming a fixed mass splitting between the gluino and lightest neutralino of 100 GeV. Limits are set for gluino lifetimes of 0.01 ns (blue), 0.1 ns (orange), 1 ns (red), and 10 ns (brown). The solid black line shows the cross section predicted by the theory. The solid colored lines show the observed upper limit for gluino lifetime, and the dashed colored lines and colored bands show the expected upper limit and its $\pm 1\sigma$ variations.

when comparing the signal efficiencies for samples with different mass splittings shown in Figures 5.17 (b) and Figures 5.18 (b).

Because the same benchmark model was used by the partial Run 2 DV + E_T^{miss} analysis [69], the exclusion limits in Figures 7.2 and 7.3 can be directly compared to the limits obtained by the previous analysis, shown in Figure 5.1. In the case of scenarios where the neutralino mass is fixed to 100 GeV, the exclusion limit has increased by 70 GeV or more for gluino lifetimes of 1 ns or less. The greatest increase in the limit occurs for gluinos with a mean proper lifetime of 0.1 ns, where the new observed limit is approximately 2550 GeV. This corresponds to an approximate increase of 190 GeV over the previous limit. More significant gains relative to the previous limits are achieved in scenarios with smaller mass splittings. Across all tested lifetimes, the new observed exclusion limit for scenarios with mass splittings of 100 GeV is at least




Figure 7.2: Exclusion limits at 95% confidence level on the mass of the gluino in the split-supersymmetry benchmark model as a function of lifetime for a fixed neutralino mass of 100 GeV. The dashed black line shows the expected limit, and the yellow band shows the area between the $\pm 1\sigma$ variations on the expected limit. The solid red contour shows the observed limit, and the dashed red lines show the observed limits when the gluino pair production cross section is varied up and down by its theoretical uncertainty. The area under the curve is excluded.

200 GeV higher than the limit established by the previous analysis. In particular, the maximum limit of approximately 2140 GeV for gluino lifetimes of 0.1 ns exceeds the limit of the previous analysis by more than 300 GeV. Regions of parameter space with $\Delta m_{\tilde{g}, \tilde{\chi}_1^0} < 100$ GeV see the most significant increase in their exclusion limits relative to the results of the previous analysis. For a gluino lifetime of 0.1 ns, a mass splitting of 30 GeV is excluded at 95% confidence level for $m_{\tilde{g}} < 1630$ GeV, shown in Figure 7.4.

Exclusion limits at 95% confidence level on the lightest neutralino mass as a function of gluino mass for a fixed gluino lifetime of 0.1 ns are shown in Figure 7.5. The region above the dashed gray line is kinematically forbidden, as the mass of the neutralino exceeds the mass of the gluino in that region. Lines parallel to the




Figure 7.3: Exclusion limits at 95% confidence level on the mass of the gluino in the split-supersymmetry benchmark model as a function of lifetime for a fixed mass splitting between the gluino and neutralino of 100 GeV. The dashed black line shows the expected limit, and the yellow band shows the area between the $\pm 1\sigma$ variations on the expected limit. The solid red contour shows the observed limit, and the dashed red lines show the observed limits when the gluino pair production cross section is varied up and down by its theoretical uncertainty. The area under the curve is excluded.

kinematically forbidden line represent constant mass splittings between the gluino and neutralino. For gluino masses less than 1600 GeV, mass splittings of 30 GeV or greater are excluded at 95% confidence level. The observed limit begins to noticeably diverge further from the kinematically forbidden line above 1800 GeV, with a maximum neutralino mass limit of approximately 2150 GeV observed at $m_{\tilde{g}} = 2350$ GeV. For higher values of $m_{\tilde{g}}$, the gain in signal efficiency from the increased mass splitting is rapidly offset by the decreasing production cross section, leading to the steep drop in the limits on $m_{\tilde{\chi}_1^0}$ seen on the right side of Figure 7.5.




Figure 7.4: Exclusion limits at 95% confidence level on the mass of the gluino in the split-supersymmetry benchmark model as a function of lifetime for a fixed mass splitting between the gluino and neutralino of 30 GeV. The dashed black line shows the expected limit, and the yellow band shows the area between the $\pm 1\sigma$ variations on the expected limit. The solid red contour shows the observed limit, and the dashed red lines show the observed limits when the gluino pair production cross section is varied up and down by its theoretical uncertainty. The area under the curve is excluded.




Figure 7.5: Exclusion limits at 95% confidence level on the mass of the neutralino in the split-supersymmetry benchmark model as a function of gluino mass for a fixed lifetime of 0.1 ns. The dashed black line shows the expected limit, and the yellow band shows the area between the $\pm 1\sigma$ variations on the expected limit. The solid red contour shows the observed limit, and the dashed red lines show the observed limits when the gluino pair production cross section is varied up and down by its theoretical uncertainty. The area under the curve is excluded.

Chapter 8

Conclusions

All right then, keep your secrets.

—Frodo Baggins

Since the discovery of the Higgs boson during Run 1 of the LHC, evidence for new physics has proved exceptionally elusive. Numerous well-motivated theories of physics beyond the Standard Model have been heavily constrained by precision measurements and searches. To minimize the possibility that new physics is hiding in the existing data set, the ATLAS Collaboration at the LHC has placed an increasing emphasis on unconventional and challenging signatures of new physics that have remained largely unexplored. There is no reason that new particles produced in collisions at the LHC must decay promptly, and there are many well-motivated models of BSM physics that predict the existence of massive, long-lived particles that could decay inside the ATLAS Inner Detector. This dissertation focuses on a search conducted by the ATLAS Collaboration for such particles in the 2016-2018 Run 2 data set.

Though the Standard Model of particle physics has proven extremely resilient to a diverse program of precision measurements, it is clear that there are questions that the SM is unable to answer. Chapter 2 outlines the SM, its successes, and its shortcomings before presenting a pair of possible supersymmetric extensions to the SM that could explain the nature of dark matter in the universe. Particular attention is given to a model of split-supersymmetry with long-lived gluinos that can decay to SM quarks and a stable neutralino, which serves as the dark matter candidate of the model. The quarks from the displaced gluino decay in the ID could create a displaced vertex with high invariant mass and high track multiplicity alongside significant E_T^{miss} from the stable neutralino LSP. Chapter 3 provides an overview of the Large Hadron

Collider and the ATLAS detector used to collect the data analyzed in this work. The reconstruction algorithms developed by the collaboration to aid in the analysis of the data are presented in Chapter 4, with particular emphasis given to the tracking and vertexing algorithms that are key to searches involving displaced vertices.

Chapter 5 provides an overview of the $DV + E_T^{\text{miss}}$ analysis, which uses the 2016-2018 Run 2 data set of 137 fb^{-1} collected at $\sqrt{s} = 13 \text{ TeV}$ to search for new, long-lived particles decaying in the ATLAS ID, leaving a signature of a displaced vertex and E_T^{miss} . This search builds on the foundation laid by previous searches carried out by the ATLAS supersymmetry group in Run 2 [2, 69, 70]. The search consists of three signal regions - one which uses the VSI secondary vertexing algorithm, and two which use the novel Fuzzy Vertexing algorithm. This dissertation focuses on the former. An extensive set of optimizations to the event selections and DV selections is performed in order to suppress the background from algorithmic and instrumental sources and maximize the signal efficiency for the benchmark split-supersymmetry model.

The data-driven methods for estimating the expected background in the signal region are detailed in Chapter 6. The primary background estimation method exploits correlations between the presence of background DVs and the track density of the event to estimate the expected number of background events in the signal region. The track density method predicts an expected background in the signal region of 0.56 ± 0.40 events. This method is successfully validated in low m_{DV} and low $N_{\text{Tracks}}^{\text{DV}}$ sideband regions. Additionally, the track density estimate is cross-checked by an alternative background estimate that predicts the contribution of individual background sources to the signal region separately in order to estimate the total background. The track density method and the combined method exhibit excellent agreement in the validation and signal regions. The remainder of Chapter 6 discusses the systematic uncertainties present in the analysis and the methods used to quantify them.

A single event is observed in the signal region, which is consistent with the predicted background. The results of the $DV + E_T^{\text{miss}}$ analysis are interpreted in the context of the benchmark split-supersymmetry model with long-lived gluinos in Chapter 7. Upper limits on the gluino pair production cross section are set as a function of the gluino mass and lifetime for scenarios with varying neutralino masses. The most stringent limits on the gluino mass are set for gluino lifetimes of 0.1 ns. In this case, gluinos are excluded up to 2550 GeV if $m_{\tilde{\chi}_1^0} = 100 \text{ GeV}$, up to 2140 GeV if $\Delta m_{\tilde{g}, \tilde{\chi}_1^0} = 100 \text{ GeV}$, and up to 1630 GeV if $\Delta m_{\tilde{g}, \tilde{\chi}_1^0} = 30 \text{ GeV}$. These limits offer significant improvements over the partial Run 2 version of the $DV + E_T^{\text{miss}}$ analysis, particularly in scenarios with

mass splittings of 100 GeV or less between the gluino and neutralino. The observed limits also compare favorably to the limits published by a recent result from the CMS Collaboration [96]. The CMS result uses the same split-supersymmetry model as a benchmark, and sets exclusion limits on the gluino mass as a function of lifetime assuming $\Delta m_{\tilde{g},\tilde{\chi}_1^0} = 100$ GeV. The CMS limit exceeds the limit on $m_{\tilde{g}}$ produced by this analysis for gluinos with lifetimes of 0.01 ns by approximately 100 GeV, but for lifetimes greater than 0.1 ns, the limit from this result exceeds the CMS limit by more than 200 GeV. For lifetimes of 1 ns or greater, the limit on $m_{\tilde{g}}$ from this work exceeds the CMS limit by 300 GeV or more.

Though this work has made significant improvements relative to previous ATLAS searches for displaced vertices, there is always room for improvement in future iterations of searches for displaced vertices. Great care was taken to minimize the background expected in the signal region, but such optimizations come at the cost of a reduction in signal efficiency. For example, many of the track cleanings used to suppress the expected background have negative effects on the signal efficiency for certain regions of parameter space. The requirement that attached tracks have 3D angle $\alpha < \pi/2$ relative to the vector pointing from the primary vertex to the DV causes a significant drop in efficiency for more massive LLPs. This occurs because more massive LLPs tend to have less significant relativistic boost, which results in a more isotropic angular distribution of their decay products. The angular track selections are found to be necessary to suppress background from accidental crossings, but alternative selections that more effectively discriminate signal DVs from background DVs are greatly desirable. The sensitivity of the analysis to compressed scenarios could also benefit from a further reduction in the offline E_T^{miss} requirement. This would require more detailed studies of the E_T^{miss} trigger efficiency in data and simulation to understand differences in the modeling of the efficiency turn-on curve.

Several improvements to the detector hardware, trigger system, and reconstruction algorithms that are being implemented for future runs of the LHC will greatly increase the sensitivity of future displaced vertex searches. Optimizations to the large-radius tracking algorithm have reduced the computational demands of the algorithm while reducing the number of fake tracks that are reconstructed. Large-radius tracking and secondary vertexing have now been implemented as part of the standard reconstruction workflow. This removes the restrictions of the DRAW_RPVLL filter and opens up the full data set to LLP analyses. New triggers that target displaced signatures using Inner Detector hits not associated to prompt tracks have also been developed for Run

3. The addition of dedicated LLP triggers and the removal of the constraints of the `DRAW_RPVLL` filter will allow for greater flexibility in the optimization of event-level selections to explore regions of phase space that previously lay out of reach.

Following Run 3, the LHC will undergo significant upgrades to provide a ten-fold increase in the integrated luminosity delivered to the experiments. The average number of pileup interactions per bunch crossing will increase from the Runs 1-3 average of 55 to approximately 200 [97]. To cope with the increased track density of such an environment, numerous upgrades to ATLAS are in development. In regards to displaced vertex searches, the most important of these upgrades is the Inner Tracker (ITk) upgrade project, which will replace the entirety of the ATLAS ID [98, 99]. The ITk will be composed entirely of silicon detectors and will expand the tracking coverage from $|\eta| < 2.5$ up to $|\eta| < 4.0$. The tracking improvements offered by the ITk and the increase in integrated luminosity delivered by the LHC will substantially increase the sensitivity of future displaced vertex searches by the ATLAS Collaboration.

Though the $DV + E_T^{\text{miss}}$ analysis detailed in this dissertation did not produce evidence of physics beyond the Standard Model, substantial improvements to existing limits on the pair production cross section of gluinos in a split-supersymmetry model have been made. Numerous models of BSM physics predict LLPs that would produce DV signatures in the ID. The three signal regions used in this search have been designed to be model-independent and provide sensitivity to a wide range of phase space in a variety of different models. Though the pair production of long-lived gluinos in a split-supersymmetry model is used as a benchmark process, the results of this search can be interpreted in the context of other models, such as the Wino-Bino coannihilation model described in Section 2.5 and a supersymmetric extension of the DFSZ axion model [100, 101]. The constraints derived from the results of this work can be used to guide future searches for new physics.

Though the elusiveness of new physics at the LHC in recent years is indeed frustrating, there is no reason to give up hope. The improvements being developed for the LHC and its experiments will open up exciting new frontiers for physicists to explore, and non-collider based experiments are continually pushing the boundaries of sensitivity to models that lie beyond the reach of current collider experiments. Though nature has refused to divulge its secrets for the time being, progress is made only through the diligent work of passionate people. The universe is a wonderful, mysterious place, and the next groundbreaking discovery could be just around the corner.

Bibliography

- [1] ATLAS Collaboration, *Search for displaced leptons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Phys. Rev. Lett.* **127** (2020) 051802, arXiv: [2011.07812](https://arxiv.org/abs/2011.07812) [[hep-ex](#)].
- [2] ATLAS Collaboration, *Search for long-lived, massive particles in events with displaced vertices and multiple jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **2023** (2023) 1, URL: [https://link.springer.com/article/10.1007/JHEP06\(2023\)200](https://link.springer.com/article/10.1007/JHEP06(2023)200).
- [3] *Standard Model of Elementary Particles*, URL: https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg (visited on 06/26/2024).
- [4] D. J. Griffiths, *Introduction to elementary particles; 2nd rev. version*, Physics textbook, New York, NY: Wiley, 2008, URL: <https://cds.cern.ch/record/111880>.
- [5] E. Noether, *Invariant variation problems*, *Transport Theory and Statistical Physics* **1** () 186, Translated from Invariante Variationsprobleme, Nachr. d. König. Gesellsch. d. Wiss. zu Göttingen, Math-phys. (1918), URL: <https://doi.org/10.1080%2F00411457108231446>.
- [6] S. Weinberg, *A Model of Leptons*, *Phys. Rev. Lett.* **19** (21 1967) 1264, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.19.1264>.
- [7] S. L. Glashow, *Partial-symmetries of weak interactions*, *Nuclear Physics* **22** (1961) 579, ISSN: 0029-5582, URL: <https://www.sciencedirect.com/science/article/pii/0029558261904692>.

- [8] P. Skands, “Introduction to QCD,” *Searching for New Physics at Small and Large Scales*, WORLD SCIENTIFIC, 2013, URL: http://dx.doi.org/10.1142/9789814525220_0008.
- [9] F. Englert and R. Brout, *Broken symmetry and the mass of gauge vector mesons*, *Physical review letters* **13** (1964) 321.
- [10] P. W. Higgs, *Broken symmetries and the masses of gauge bosons*, *Physical review letters* **13** (1964) 508.
- [11] 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](https://arxiv.org/abs/1207.7214) [[hep-ex](#)].
- [12] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Physics Letters B* **716** (2012) 30, arXiv: [1207.7235](https://arxiv.org/abs/1207.7235) [[hep-ex](#)].
- [13] L. Lee, C. Ohm, A. Soffer, and T.-T. Yu, *Collider searches for long-lived particles beyond the Standard Model*, *Progress in Particle and Nuclear Physics* **106** (2019) 210, ISSN: 0146-6410, URL: <https://www.sciencedirect.com/science/article/pii/S0146641019300109>.
- [14] UA1 Collaboration, *Experimental observation of events with large missing transverse energy accompanied by a jet or a photon (S) in pp collisions at $s = 540$ GeV*, *Physics Letters B* **139** (1984) 115, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/0370269384900467>.
- [15] UA2 Collaboration, *Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider*, *Physics Letters B* **122** (1983) 476, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/0370269383916052>.
- [16] UA1 Collaboration, *Experimental observation of lepton pairs of invariant mass around 95 GeV/c² at the CERN SPS collider*, *Physics Letters B* **126** (1983) 398, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/0370269383901880>.

- [17] UA2 Collaboration, *Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN pp collider*, *Physics Letters B* **129** (1983) 130, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/037026938390744X>.
- [18] CDF Collaboration, *Observation of Top Quark Production in $P\bar{p}$ -P Collisions with the Collider Detector at Fermilab*, *Physical Review Letters* **74** (1995), ISSN: 1079-7114, URL: <http://dx.doi.org/10.1103/PhysRevLett.74.2626>.
- [19] D0 Collaboration, *Observation of the Top Quark*, *Physical Review Letters* **74** (1995), ISSN: 1079-7114, URL: <http://dx.doi.org/10.1103/PhysRevLett.74.2632>.
- [20] J. J. Aubert et al., *Experimental Observation of a Heavy Particle J*, *Phys. Rev. Lett.* **33** (23 1974) 1404, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.33.1404>.
- [21] J. E. Augustin et al., *Discovery of a Narrow Resonance in e^+e^- Annihilation*, *Phys. Rev. Lett.* **33** (23 1974) 1406, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.33.1406>.
- [22] ATLAS Collaboration, *Combined measurement of the Higgs boson mass from the $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ^* \rightarrow 4\ell$ decay channels with the ATLAS detector using $\sqrt{s} = 7, 8$ and 13 TeV pp collision data*, *Phys. Rev. Lett.* **131** (2023) 251802, arXiv: 2308.04775, URL: <https://cds.cern.ch/record/2867152>.
- [23] L. Bergström, *Non-baryonic dark matter: observational evidence and detection methods*, *Reports on Progress in Physics* **63** (2000), ISSN: 1361-6633, URL: <http://dx.doi.org/10.1088/0034-4885/63/5/2r3>.
- [24] Particle Data Group, *Review of Particle Physics*, *PTEP* **2022** (2022) 083C01, URL: <https://academic.oup.com/ptep/article/2022/8/083C01/6651666>.
- [25] S. P. MARTIN, “A SUPERSYMMETRY PRIMER,” *Advanced Series on Directions in High Energy Physics*, WORLD SCIENTIFIC, 1998, URL: http://dx.doi.org/10.1142/9789812839657_0001.

- [26] G. Arcadi, A. Djouadi, and M. Kado, *The Higgs-portal for dark matter: effective field theories versus concrete realizations*,
The European Physical Journal C **81** (2021) 1, URL: <https://link.springer.com/article/10.1140/epjc/s10052-021-09411-2>.
- [27] M. J. Strassler and K. M. Zurek, *Echoes of a hidden valley at hadron colliders*,
Physics Letters B **651** (2007), ISSN: 0370-2693,
URL: <http://dx.doi.org/10.1016/j.physletb.2007.06.055>.
- [28] A. Arvanitaki, N. Craig, S. Dimopoulos, and G. Villadoro, *Mini-Split*,
Journal of High Energy Physics **2013** (2013), ISSN: 1029-8479,
URL: [http://dx.doi.org/10.1007/JHEP02\(2013\)126](http://dx.doi.org/10.1007/JHEP02(2013)126).
- [29] A. Arvanitaki, C. Davis, P. W. Graham, A. Pierce, and J. G. Wacker,
Limits on split supersymmetry from gluino cosmology,
Physical Review D **72** (2005), ISSN: 1550-2368,
URL: <http://dx.doi.org/10.1103/PhysRevD.72.075011>.
- [30] K. Rolbiecki and K. Sakurai,
Long-lived bino and wino in supersymmetry with heavy scalars and higgsinos,
JHEP **2015** (2015) 1,
URL: [https://link.springer.com/article/10.1007/JHEP11\(2015\)091](https://link.springer.com/article/10.1007/JHEP11(2015)091).
- [31] *CERN: Our Member States*,
<https://home.cern/about/who-we-are/our-governance/member-states>,
Accessed: 2024-4-26.
- [32] L. Evans and P. Bryant, *LHC Machine*,
Journal of Instrumentation **3** (2008) S08001,
URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08001>.
- [33] E. A. Mobs,
The CERN accelerator complex. Complexe des accélérateurs du CERN,
(2016), General Photo, URL: <https://cds.cern.ch/record/2225847>.
- [34] R. D. Ball et al.,
Parton distributions from high-precision collider data: NNPDF Collaboration,
The European Physical Journal C **77** (2017), ISSN: 1434-6052,
URL: <http://dx.doi.org/10.1140/epjc/s10052-017-5199-5>.

- [35] T. Gleisberg et al., *Event generation with SHERPA 1.1*, *Journal of High Energy Physics* **2009** (2009) 007, ISSN: 1029-8479, URL: <http://dx.doi.org/10.1088/1126-6708/2009/02/007>.
- [36] ATLAS Collaboration, *Number of Interactions per Crossing*, CERN, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2> (visited on 04/29/2024).
- [37] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, *The European Physical Journal C* **83** (2023), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-023-11747-w>.
- [38] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003.
- [39] J. Pequeno and P. Schaffner, *How ATLAS Detects Particles: Diagram of Particle Paths in the Detector*, <https://cds.cern.ch/record/1505342>, 2013.
- [40] *ATLAS Magnet System Diagram*, <http://www.quantumdiaries.org/wp-content/uploads/2011/05/exp-magnets.png>, Accessed: 2019-5-23.
- [41] ATLAS Collaboration, *ATLAS insertable B-layer technical design report*, CERN-LHCC-2010-013 (2010), URL: <https://cds.cern.ch/record/1291633>.
- [42] ATLAS Collaboration, *Efficiency and Hit Spatial Resolution of ATLAS IBL Sensors in LHC Run 2 Collision Events*, ATL-INDET-PUB-2016-001, 2016, URL: <https://cds.cern.ch/record/2203893>.
- [43] ATLAS Collaboration, *Study of the material of the ATLAS inner detector for Run 2 of the LHC*, *Journal of Instrumentation* **12** (2017), ISSN: 1748-0221, URL: <http://dx.doi.org/10.1088/1748-0221/12/12/P12009>.
- [44] K. Potamianos, *The upgraded Pixel detector and the commissioning of the Inner Detector tracking of the ATLAS experiment for Run-2 at the Large Hadron Collider*, 2016, arXiv: [1608.07850](https://arxiv.org/abs/1608.07850) [physics.ins-det], URL: <https://arxiv.org/abs/1608.07850>.

- [45] J. Pequeno, “Computer Generated image of the ATLAS calorimeter,” 2008, URL: <https://cds.cern.ch/record/1095927>.
- [46] *ATLAS muon spectrometer: Technical Design Report*, Technical design report. ATLAS, Geneva: CERN, 1997, URL: <https://cds.cern.ch/record/331068>.
- [47] ATLAS Collaboration, *Operation of the ATLAS trigger system in Run 2*, *JINST* **15** (2020) P10004, arXiv: 2007.12539 [hep-ex].
- [48] ATLAS Collaboration, *Performance of the ATLAS trigger system in 2015*, *Eur. Phys. J. C* **77** (2017) 317, arXiv: 1611.09661 [hep-ex].
- [49] M. Elsing, L. Goossens, A. Nairz, and G. Negri, *The ATLAS Tier-0: Overview and operational experience*, *Journal of Physics: Conference Series* **219** (2010) 072011, URL: <https://dx.doi.org/10.1088/1742-6596/219/7/072011>.
- [50] A. Salzburger, *Inner Detector Reconstruction Part 1 - Tracking*, 2008, URL: https://indico.cern.ch/event/34087/contributions/802944/attachments/669263/919934/IDTracking_A_Salzburger.pdf (visited on 05/30/2024).
- [51] ATLAS Collaboration, *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*, *Eur. Phys. J. C* **77** (2017) 673, arXiv: 1704.07983, URL: <https://cds.cern.ch/record/2261156>.
- [52] R. Frühwirth, *Application of Kalman filtering to track and vertex fitting*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **262** (1987) 444, ISSN: 0168-9002, URL: <https://www.sciencedirect.com/science/article/pii/0168900287908874>.
- [53] T. Cornelissen et al., *The new ATLAS track reconstruction (NEWT)*, *Journal of Physics: Conference Series* **119** (2008) 032014, URL: <https://dx.doi.org/10.1088/1742-6596/119/3/032014>.
- [54] ATLAS Collaboration, *Performance of the reconstruction of large impact parameter tracks in the inner detector of ATLAS*, ATL-PHYS-PUB-2017-014, 2017, URL: <https://cds.cern.ch/record/2275635>.

- [55] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, *Eur. Phys. J. C* **77** (2017) 332, arXiv: 1611.10235 [hep-ex].
- [56] ATLAS Collaboration, *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*, ATL-PHYS-PUB-2019-013, 2019, URL: <https://cds.cern.ch/record/2669425>.
- [57] H. Oide, *VrtSecInclusive: an inclusive displaced vertex reconstruction package in release 21*, tech. rep., CERN, 2018, URL: <https://cds.cern.ch/record/2303549>.
- [58] D. Rousso, *Searching for long-lived particles Beyond the Standard Model with displaced vertex signatures using the ATLAS detector: Not all hopes for new physics die immediately*, PhD thesis: University of Cambridge, 2023, URL: <https://cds.cern.ch/record/2876819>.
- [59] ATLAS Collaboration, *Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **81** (2020) 689, arXiv: 2007.02645 [hep-ex].
- [60] M. Cacciari, G. P. Salam, and G. Soyez, *The anti- k_t jet clustering algorithm*, *Journal of High Energy Physics* **2008** (2008) 063, URL: <https://iopscience.iop.org/article/10.1088/1126-6708/2008/04/063>.
- [61] 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].
- [62] 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].
- [63] ATLAS Collaboration, *Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton–proton collision data*, *JINST* **14** (2019) P03017, arXiv: 1812.03848 [hep-ex].
- [64] 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].

- [65] M. Cacciari, G. P. Salam, and G. Soyez, *The catchment area of jets*, *Journal of High Energy Physics* **2008** (2008) 005, ISSN: 1029-8479, URL: <http://dx.doi.org/10.1088/1126-6708/2008/04/005>.
- [66] ATLAS Collaboration, *Performance of algorithms that reconstruct missing transverse momentum in $\sqrt{s} = 8$ TeV proton–proton collisions in the ATLAS detector*, *Eur. Phys. J. C* **77** (2017) 241, arXiv: 1609.09324 [hep-ex].
- [67] ATLAS Collaboration, *Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13$ TeV*, *The European Physical Journal C* **78** (2018) 1, URL: <https://link.springer.com/article/10.1140/epjc/s10052-018-6288-9>.
- [68] ATLAS Collaboration, *Performance of the missing transverse momentum triggers for the ATLAS detector during Run-2 data-taking*, *JHEP* **08** (2020) 080, arXiv: 2005.09554 [hep-ex].
- [69] ATLAS Collaboration, *Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Phys. Rev. D* **97** (2018) 052012, arXiv: 1710.04901 [hep-ex].
- [70] ATLAS Collaboration, *Search for long-lived, massive particles in events with a displaced vertex and a muon with large impact parameter in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **102** (2020) 032006, arXiv: 2003.11956 [hep-ex].
- [71] 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 [hep-ex].
- [72] R. Brun and F. Rademakers, *ROOT – An object oriented data analysis framework*, *Nucl. Instrum. Meth. A* **389** (1997) 81, ISSN: 0168-9002.
- [73] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079, arXiv: 1405.0301 [hep-ph].

- [74] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159, arXiv: 1410.3012 [hep-ph].
- [75] R. D. Ball et al., *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244, arXiv: 1207.1303 [hep-ph].
- [76] GEANT4 Collaboration, S. Agostinelli, et al., *GEANT4 – a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250.
- [77] E. Bothmann et al., *Event generation with Sherpa 2.2*, *SciPost Phys.* **7** (2019) 034, arXiv: 1905.09127 [hep-ph].
- [78] ATLAS Collaboration, *Characterisation and mitigation of beam-induced backgrounds observed in the ATLAS detector during the 2011 proton–proton run*, *JINST* **8** (2013) P07004, arXiv: 1303.0223 [hep-ex].
- [79] ATLAS Collaboration, *Selection of jets produced in 13 TeV proton–proton collisions with the ATLAS detector*, ATLAS-CONF-2015-029, 2015, URL: <https://cds.cern.ch/record/2037702>.
- [80] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 639, arXiv: 1902.04655 [hep-ex].
- [81] ATLAS Collaboration, *Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016*, *Eur. Phys. J. C* **79** (2019) 205, arXiv: 1810.05087 [hep-ex].
- [82] M. Oreglia, *A Study of the Reactions $\psi' \rightarrow \gamma\gamma\psi$* , MSc thesis: Stanford University, 1980.
- [83] ATLAS Collaboration, *Early Inner Detector Tracking Performance in the 2015 Data at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2015-051, 2015, URL: <https://cds.cern.ch/record/2110140>.
- [84] W. Beenakker, R. Höpker, M. Spira, and P. Zerwas, *Squark and gluino production at hadron colliders*, *Nucl. Phys. B* **492** (1997) 51, arXiv: hep-ph/9610490.

- [85] A. Kulesza and L. Motyka, *Threshold Resummation for Squark-Antisquark and Gluino-Pair Production at the LHC*, *Phys. Rev. Lett.* **102** (2009) 111802, arXiv: [0807.2405](https://arxiv.org/abs/0807.2405) [[hep-ph](#)].
- [86] A. Kulesza and L. Motyka, *Soft gluon resummation for the production of gluino-gluino and squark-antisquark pairs at the LHC*, *Phys. Rev. D* **80** (2009) 095004, arXiv: [0905.4749](https://arxiv.org/abs/0905.4749) [[hep-ph](#)].
- [87] W. Beenakker et al., *Soft-gluon resummation for squark and gluino hadroproduction*, *JHEP* **12** (2009) 041, arXiv: [0909.4418](https://arxiv.org/abs/0909.4418) [[hep-ph](#)].
- [88] W. Beenakker et al., *NNLL resummation for squark-antisquark pair production at the LHC*, *JHEP* **01** (2012) 076, arXiv: [1110.2446](https://arxiv.org/abs/1110.2446) [[hep-ph](#)].
- [89] W. Beenakker et al., *Towards NNLL resummation: hard matching coefficients for squark and gluino hadroproduction*, *JHEP* **10** (2013) 120, arXiv: [1304.6354](https://arxiv.org/abs/1304.6354) [[hep-ph](#)].
- [90] W. Beenakker et al., *NNLL resummation for squark and gluino production at the LHC*, *JHEP* **12** (2014) 023, arXiv: [1404.3134](https://arxiv.org/abs/1404.3134) [[hep-ph](#)].
- [91] J. Butterworth et al., *PDF₄LHC recommendations for LHC Run II*, *Journal of Physics G: Nuclear and Particle Physics* **43** (2016) 023001, ISSN: 1361-6471, URL: <http://dx.doi.org/10.1088/0954-3899/43/2/023001>.
- [92] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *The European Physical Journal C* **71** (2011), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-011-1554-0>.
- [93] A. L. Read, *Presentation of search results: the CL_S technique*, *J. Phys. G* **28** (2002) 2693.
- [94] L. Heinrich, M. Feickert, G. Stark, and K. Cranmer, *pyhf: pure-Python implementation of HistFactory statistical models*, *J. Open Source Softw.* **6** (2021) 2823.

- [95] L. Heinrich, M. Feickert, and G. Stark, *pyhf: v0.7.6*, version 0.7.6, <https://github.com/scikit-hep/pyhf/releases/tag/v0.7.6>, URL: <https://doi.org/10.5281/zenodo.1169739>.
- [96] CMS Collaboration, *Search for long-lived particles using displaced vertices and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Physical Review D* **109** (2024), ISSN: 2470-0029, URL: <http://dx.doi.org/10.1103/PhysRevD.109.112005>.
- [97] G. Apollinari et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report V. 0.1*, CERN Yellow Reports: Monographs, Geneva: CERN, 2017, URL: <https://cds.cern.ch/record/2284929>.
- [98] ATLAS Collaboration, *ATLAS Inner Tracker Strip Detector: Technical Design Report*, ATLAS-TDR-025; CERN-LHCC-2017-005, 2017, URL: <https://cds.cern.ch/record/2257755>.
- [99] ATLAS Collaboration, *ATLAS Inner Tracker Pixel Detector: Technical Design Report*, ATLAS-TDR-030; CERN-LHCC-2017-021, 2017, URL: <https://cds.cern.ch/record/2285585>.
- [100] M. Dine, W. Fischler, and M. Srednicki, *A simple solution to the strong CP problem with a harmless axion*, *Phys. Lett. B* **104** (1981) 199, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/0370269381905906>.
- [101] G. Hoshino, *Searching for the DFSZ Axino in Collider Experiments*, Presented at *PHENO 2023*, Pittsburgh, 2023, URL: <https://indico.cern.ch/event/1218225/contributions/5384880/>.