

**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**

David Rousso
Churchill College
University of Cambridge



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

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

David Rousso

Traditional LHC searches assume that Beyond the Standard Model particles decay immediately after they are produced. However, there is no a priori motivation for this assumption. This thesis details part of the growing effort to use new unorthodox signatures to search for models where these particles are long-lived. All searches were conducted with the full 139 fb^{-1} of Run 2 data taken in 2015-2018.

The first analysis uses displaced vertex signatures in multijet-triggered events to search for long-lived particles that decay into hadrons in the ATLAS inner detector. This was the first such Run 2 search, placing the best limits in this lifetime range to date on R-parity violating supersymmetry models. The second analysis uses a similar strategy, but using missing transverse energy as the trigger. This improves upon the early Run 2 2016-data search with an updated analysis strategy and a new vertexing algorithm. A focus on interpreting this search with the the Higgs Portal model predicts sensitivity to the exclusion of the Higgs branching ratio to hadronically-decaying long-lived particles to the lowest limit to date.

As an outlook to the future of displaced vertex analyses using the ATLAS inner tracker upgrade (ITk), a script library is also described for the automation of quality control for its strip sensors.

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text.

It is not substantially the same as any work that has already been submitted before for any degree or other qualification except as declared in the preface and specified in the text.

It does not exceed the prescribed word limit for the Faculty of Physics & Chemistry Degree Committee.

David Rousso

Acknowledgements

Despite the common physics assumptions, a PhD does not occur in a vacuum. Particularly for a PhD that took place during a global pandemic and somehow continuously unprecedented times, there are countless people without which this PhD would have been significantly less tolerable or even impossible. This is by far not a complete list, and I apologize profusely to anybody that I may have missed.

Perhaps most important to acknowledge are my supervisors Tina and Oleg, whose powers combined have resulted in the best super-supervisor I could have ever wished for. Their complementary strengths have moulded me from a nanotechnology engineering student who knew absolutely nothing about particle physics into a proper researcher in ATLAS with my own niche of expertise. There were many times during the PhD where I had felt hopelessly stressed and didn't know what to do, thinking that this wasn't for me and wanting to throw in the towel. Despite entering meetings anxious and unsure of everything, Tina had a magical ability to, even without telling her about my worries, making it such that I would leave every meeting clear-headed, ready to conquer the week, and wondering what I was ever worrying about. On the days where I felt stuck and with stagnant ambitions, I always knew I could rely on Oleg and his directness to give me advice on how to improve and grow as a researcher, keep me on track, and keep my critical thinking skills on my toes by testing my knowledge and making sure I fully understood what I was doing. You are my role models, and if I could one day have even a third of my life together as you both do when I "grow up" I would be incredibly happy with myself. You make an amazing team, and I hope that you continue to guide many future students the same way you guided me.

Similarly, I owe my qualification task local supervisor Bart my gratitude for being amazing in training me in the cleanroom and quality control tests, as well as his infinite patience and enthusiasm with our many meetings together as I tried to hash out the design of the QC scripts. Similarly I owe my gratitude to Vitaliy for being my

CERN-side supervisor for the qualification task, for helping me integrate into the ITk collaboration, and helping make the QC scripts successful.

To the DV+Jets team, particularly Giulia and Emily, thank you for putting up with me and working with me with various parts of the analysis during the pandemic. You are both amazing physicists, and I hope I can follow in your footsteps. To the DV+MET team, thank you for being an amazing and fun team to work with, especially with all the memes. In particular, thank you to Evan, my co-analysis contact. It has been an honour driving this flaming bus off a cliff together.

I am also greatly indebted to Will, Jon, and Holly, our postdocs who have always had their door open to when I needed advice or needed to talk. I would also like to thank the Cambridge ATLAS Analysis Anonymous ("AA") 2019 Group (Dom and Ben), Fionn, as well as everyone else in our lab for making it such a fun environment. To the AA in particular, it's been great sharing offices with you guys and commiserating about the failed parts of our PhDs. Dom has been an amazing partner in crime with the QC scripts, and I have enjoyed our adventures together dealing with the ITk unicorns. I also owe a special thanks to Lars for always having answers to my many questions while we shared an office. To Anna, it has been an absolute pleasure working with you and guiding you, and thank you for putting up with me as I continue to learn how to guide people. As well as Anna I would also like to thank Seb and Richard for making the post-exodus part of LTA so much fun. To Jeremy, it has been an absolute pleasure going from supervising and coxing you to seeing you join our group. I will always cherish our physics discussions. I also must thank Steve and Julia for being amazing and helping make sure everything on the computation side runs smoothly.

To my close friends at Cambridge Gates: GARBO (William and Michelle) and Ellen, thank you for making my time at Cambridge so special. It has been amazing fun having great adventures with you all, hot chocolates, good food, museums, and general memery. Thank you for making my return to Cambridge so seamless. I also need to thank Mika for our many walks together during the pandemic that have made such chaotic times very calming. All of you inspire me every day with your incredibly cool research, and it has been an honour to become close to such amazing people as you.

It was an absolute privilege to have become friends with so many amazing people at CERN during my LTA. Ynyr, Anja, Neza, Iza, Riz, Jenny, Sukanya, Benjamin, Emir, and so many others who have inspired me every day to be the best physicist I can be.

Thank you for our interesting physics discussions together, not to mention making LTA incredibly fun. Despite difficult times, you always make me remember how amazing and exciting things are in our field. Thank you for making my time at CERN so special. I hope we can work together on a cool analysis in the future.

My many housemates over the years also deserve my eternal gratitude, particularly for making the pandemic as pleasurable as possible. In particular a special thanks to Jesse for having taught me how to supervise, Jye and Dao for amazing rice-cooker hot-pot, and Chris for being my go-to theorist who also introduced me to coxing. I would also like to express my gratitude to Product River Corp. (Sam and Nelson), with whom our weekly scrum meetings during the pandemic had truly made it a special and much more fun experience.

A special thanks is in order to the UK (Ashlyn, Austin, and Ethan) and Swiss (Manasa, Mary, Jeremie, and Theeran) NanoHUBs for all of the fun adventures and travels. You have made my country moves feel much less chaotic by always knowing I'll have good friends on either side. Your ability to manifest the most random suggestions in conversation into full-fledged reality has taught me that one mustn't just let dreams be dreams. I would like to thank Manasa ("David Rousso") in particular for always offering amazing and actionable advice, for holding me accountable for my dreams, and giving me opportunities to reach outside of particle physics.

I of course need to thank all of my previous supervisors and their PIs, who have helped shape my career path to get me here today. Reza (and Federico) led my first ever introduction to research and has greatly benefitted my career by planning out projects such that I would get papers. Ida (and Joanna) thank you for training me in the cleanroom and going out of your way to make sure I acquired such important skills that would benefit the rest of my career. Rapha (and Rob and Zoran), although it may not have been intentional, have taught me about European culture, lasers, and the important lesson that although one may look dumb for a minute while asking a dumb question, if one does not ask that question they remain dumb for the rest of their life. Yamagishi-san, Shinozuka-san, and Ono-san (and Chris) have taught me the art of asking useful and objective-oriented questions, design, not having fear of just going for things, and analysing network traffic when API documentation is non-existent. Anna (and Andreas and Aldo and Klaus), thanks for having taken a chance on me, introduced me to the world of particle physics, helped me fall in love with the subject, and taught me about the wonders of Switzerland.

My family deserves enormous thanks, who, spread out across the world, has supported me during my many country moves on my path here. In particular, I would like to thank my aunt Alice, the Suens, and the Betts for taking me in and helping me feel at home in the UK. I would also like to thank my uncle Dr. Wong who had inspired me to do outreach and has been an inspiring role model since I was a kid. I feel honoured to be soon joining your ranks.

And finally, in the words of Snoop Dogg, "I want to thank me", for having persevered all those years and believed in myself for long enough to make it here.

Preface

The Standard Model describes the fundamental constituents of the universe and their interactions with each other. Although it has been remarkably successful in predicting observed phenomena, there are still some large open questions, such as the Higgs mass problem and dark matter, which strongly hint that there lie new as-of-yet undiscovered particles Beyond the Standard Model. Most searches in the ATLAS detector search for such new particles with the assumption that either they either decay immediately after they are produced in the proton-proton collision at its centre, or that they are stable neutral particles and are completely invisible to the detector. In the former case, only its decay products get detected by the detector itself. In the latter case, the only evidence that remains is a missing transverse momentum. However, there is no a priori motivation for why these new particles should only either decay immediately or be stable. Particles that live just long enough to travel into the detector before decaying can result in interesting and new signatures that need bespoke search strategies. It is imperative that the possibility that these particles are "long-lived" is checked. The signature focused on in this work is the displaced vertex, when the long-lived particle travels a few cm from the interaction point and then decays into many charged tracks in the ATLAS inner detector.

Chapter 1 focuses on elucidating the motivation behind this work. This begins by a summary of the current Standard Model of particle physics, before going into some of its shortfalls that are relevant to this work which hint at the existence of new physics Beyond the Standard Model. The motivation for searching for such new physics in the form of long-lived particles is then outlined, as well as a couple of relevant benchmark models used in the thesis for analysis design and interpretation.

Chapter 2 describes the experimental apparatus used in this thesis, the ATLAS detector and Large Hadron Collider at CERN in Geneva. Focus will be given to the Inner Detector in particular due to its importance to the displaced vertex signature that is central to this work. This will also describe important physics analysis objects

such as jets and missing transverse energy, as well as describe the existing process of reconstructing displaced tracks and grouping them together to create the displaced vertex signature.

Both searches described in this work, the DV+Jets search in Chapter 3 and the DV+MET search in Chapter 4, look for Beyond the Standard Model particles in the case that they may be long-lived and decay into many charged particles such that they produce displaced vertex signatures in the Inner Detector of ATLAS for the full Run 2 data-taking from 2015-2018.

The first search is the first search of its kind for Run 2. The author was part of the relatively small core team in the analysis amongst a larger team of ATLAS colleagues. The results were published [1] and the author presented the pre-published public results at the EDSU2022 conference in November 2022, as well as along other LHC long-lived particle results at the BLV2022 conference in September 2022.

The second search extends the sensitivity of an earlier result made with only early Run 2 data from 2016 [2]. The analysis incorporates many improvements, learning lessons from the DV+Jets analysis to create a more robust analysis design. Additionally, the new "fuzzy" vertexing algorithm developed by Risa Ushioda allows this analysis to dramatically increase its sensitivity to the case where the LLP decays into heavy quarks. Hadrons with heavy quarks have a slightly long lifetime of their own, which result in effectively sub-vertices that are slightly displaced from the main LLP decay displaced vertex. The author was one of the standard two appointed "Analysis Contacts" of the analysis, being the leader of the analysis. Working with a team of other ATLAS colleagues, the author took a leading role in designing the analysis and was heavily involved in most aspects of the analysis through the guiding of other students. The author's substantive work was primarily on redeveloping the material map veto and all aspects relating to a particular signal region requiring two displaced fuzzy vertices, which has been designed using the Higgs Portal model as the benchmark model. Here, the goal is to obtain the highest sensitivity to date for limits on the branching ratio of the Higgs to a hadronically-decaying long-lived particle pair. The various aspects of Monte Carlo production needed to produce this result was split between the author and Anna Mullin. The final interpretation was done by the author.

Chapter 5 describes the author's work in the ATLAS inner tracker upgrade (ITk), carried out to become a full signing member of the ATLAS Collaboration. This outlines the motivation and then the details of the task of developing a software

framework for the automatic parameter extraction, evaluation, database interfacing, and reporting of quality control tests on strip sensors. Due to the unusually large scope of the qualification task, part of this task was attached to that of Dominic Jones, who worked on a few of the QC test evaluation modules within the framework as well as helped with maintaining the repository. Work on this task was presented along with proceedings at the 2021 Pisa Meeting held in May 2022 [3].

Finally, Chapter 6 contains concluding remarks.

Contents

1	Theoretical Background and Motivation	1
1.1	The Standard Model	1
1.2	The Motivation Behind New Physics	5
1.2.1	The Higgs Mass Problem	5
1.2.2	Dark Matter	7
1.3	The Motivation for Long Lived Particles	8
1.4	Selected Benchmark Models	11
1.4.1	Higgs/Scalar Portal	11
1.4.2	Supersymmetry (SUSY)	14
2	Experimental Background: The LHC and the ATLAS Detector	17
2.1	The LHC at CERN	17
2.2	The ATLAS Experiment	20
2.2.1	The Inner Detector (ID)	22
2.2.2	The Electromagnetic and Hadronic Calorimeters (ECAL and HCAL)	25
2.2.3	The Muon Spectrometer (MS)	26
2.3	Object Reconstruction	27
2.3.1	Jets	27
2.3.2	Missing Transverse Energy (MET)	29
2.3.3	Standard Tracking	31
2.3.4	Large-Radius Tracking and Displaced Vertex Reconstruction	31
3	The Displaced Vertex + Jets Search (DV+Jets)	35
3.1	Data and MC Samples	38
3.2	Selections and Signal Regions	39
3.2.1	Event Selections	40
3.2.2	Displaced Vertex Selections	43
3.2.3	Material Map Veto	45

3.2.4	Track Cleanings	48
3.3	Sources of DV Backgrounds	51
3.3.1	Framework for Truth Studies	51
3.3.2	General Breakdown of DV Backgrounds	60
3.3.3	Truth Single-Process DVs (tSP)	62
3.3.4	Truth Single-Process Plus Track DVs (tAX)	65
3.3.5	Other Categories of Truth DVs (tOther)	68
3.4	The Combined Background Estimate	71
3.4.1	Estimating Single-Process DVs (eSP)	71
3.4.2	Estimating Accidental Crossings (eAX)	78
3.4.3	Estimating Merged Vertices (eMV)	84
3.4.4	Closure Uncertainty for Accidental Crossings (eAX) and Merged Vertices (eMV)	85
3.4.5	Combining Background Estimates	86
3.5	Inclusive Background Estimates	87
3.6	Uncertainties	89
3.7	Results	90
4	The Displaced Vertex + Missing Transverse Energy Search (DV+MET)	95
4.1	Samples and Basic Selections	99
4.2	Updating the Material Map Veto	101
4.3	Determining the Signal Regions	112
4.3.1	Fuzzy DV Signal Regions Selections	112
4.3.2	VSI DV Signal Region Selections	121
4.4	Higgs Portal Sensitivity Study	125
4.5	Conclusion and Future Steps	127
5	ATLAS Inner Tracker Upgrade (ITk) Strips Sensors Quality Control	129
5.1	Workflow	132
5.2	An Example of QC Analysis: Treatment of CV Tests	133
5.3	Batch Reporting	135
5.4	Current Status	136
6	Conclusion and Outlook	137
	List of figures	153
	List of tables	161

“There are no plans, only vibes”

— Manasa Kani, PhD Computational Nanoelectronics, ETH Zurich

“I can see 2 m ahead, but unfortunately I’m going at 5 m/s...”

— Ynyr Harris, PhD Particle Physics, ATLAS, University of Oxford

Chapter 1

Theoretical Background and Motivation

“Well, it’s not the worst argument theorists have made...”

— Lars Henkelmann, PhD Particle Physics, ATLAS, University of Cambridge

1.1 The Standard Model

The Standard Model (SM) of particle physics is the current established theory for the description of interactions of fundamental particles in high energy physics. The experimental confirmation of the existence of all its current particles was completed in 2012 with the discovery of the Higgs boson at the LHC [4, 5]. Although the SM is now well-established, there are still some phenomena involving known particles it does not incorporate, such as explaining neutrino oscillations [6], and there may be extensions in the future.

The need to make quantum mechanics relativistic led to the development of the first quantum field theory, quantum electrodynamics (QED), which described the interactions of electrically charged spin-1/2 fermions and photons, where photons are the particles that mediate the electromagnetic force. This also required the proposal and discovery of antiparticles, where it was found that every fermion should have a corresponding anti-fermion. After the discovery was made that the proton was actually made of three quarks held together by gluons, quantum chromo-dynamics (QCD) was developed to additionally explain the strong force and the colour charge.

After the weak force was discovered via the observation of charge-parity violating beta-decay, along with discoveries of neutrino oscillations, the model was also updated to describe these. However, it was discovered that with this new model, symmetry would not allow particles to have an intrinsic mass, and that in order for them to have this, it must be effectively imparted to them via an interaction with a scalar field, the Higgs field. With this final addition, we arrive at the SM that we know today [7].

In general, the SM comprises fundamental quantum fields which can interact with each other in certain ways, where we can observe excitations of these fields as the elementary particles. This quantum field theory describes 12 fermions (6 quarks, 3 neutrinos, and 3 charged leptons), which comprise what we typically think of as matter, and all have corresponding anti-fermions. The SM also describes 4 classes of gauge bosons which we typically think of as force carriers, the 8 gluons g of the strong force, the photon γ of the electromagnetic force, and the W^\pm and Z bosons of the weak force. Finally, we have one scalar boson, the Higgs, which allows most elementary particles to have mass by interacting with them [8, 9]. A summary of the SM particles and their properties is outlined in Figure 1.1.

Spin ½ Fermions				Spin 1 Boson	Spin 0 Boson (Scalar)
Quarks	$Q = \frac{2}{3}, I_{3,L} = \frac{1}{2}$	<div>2.16 MeV</div> <div>u</div> <div>up</div>	<div>1.27 GeV</div> <div>c</div> <div>charm</div>	<div>173 GeV</div> <div>t</div> <div>top</div>	<div>(massless)</div> <div>γ</div> <div>photon</div>
	$Q = -\frac{1}{3}, I_{3,L} = -\frac{1}{2}$	<div>4.67 MeV</div> <div>d</div> <div>down</div>	<div>93.4 MeV</div> <div>s</div> <div>strange</div>	<div>4.18 GeV</div> <div>b</div> <div>bottom</div>	<div>(massless)</div> <div>g</div> <div>gluon</div>
Leptons	$Q = -1, I_{3,L} = -\frac{1}{2}$	<div>0.511 MeV</div> <div>e^-</div> <div>electron</div>	<div>106 MeV</div> <div>μ^-</div> <div>muon</div>	<div>1.78 GeV</div> <div>τ^-</div> <div>tau</div>	<div>80.4 GeV</div> <div>W^\pm</div> <div>W boson</div>
	$Q = 0, I_{3,L} = \frac{1}{2}$ (neutrinos)	<div><1.1 eV</div> <div>ν_e</div> <div>electron neutrino</div>	<div><0.19 MeV</div> <div>ν_μ</div> <div>muon neutrino</div>	<div><18.2 MeV</div> <div>ν_τ</div> <div>tau neutrino</div>	<div>91.2 GeV</div> <div>Z</div> <div>Z boson</div>

Figure 1.1: A summary of the fundamental particles of the SM. The fermions have anti-fermion counterparts. Particle masses from [10].

The quarks, which interact with the strong force via the gluons, possess a strong "charge" known as a colour charge. There are 3 such charges: red, green, and blue for

quarks and anti-red, anti-green, and anti-blue for anti-quarks. There is nothing special about any of the individual colours, and therefore the labels are arbitrary. Therefore, the physics should be invariant about the exchange of any of the colours. This means that QCD has symmetry about the exchange of the 3 colours. This symmetry is referred to mathematically as QCD belonging to the $SU(3)_{\text{colour}}$ symmetry group, which also represents the conservation of colour charge through Noether's theorem [11]. The strong force is very strong compared to the other forces, and due to the gluon's ability to self-interact, this results in the strong potential with respect to distance increasing roughly linearly rather than decreasing with an inverse square like in electromagnetics. This results in quarks not being observable on their own. Instead, only observable in colourless bound states called hadrons. When quarks are involved in high energy processes that would cause them to be separated, for example just being produced in a collision at the LHC and travelling in opposite directions, it is energetically favourable for quark-antiquark pairs to be produced such that every quark ends up in a bound hadron colourless state. This results in a shower of hadrons being emitted in the direction of each original quarks, known as jets.

The weak force only interacts with fermions with a left-handed chirality. Unlike helicity (spin with respect to momentum direction) which can change depending on frame of reference, chirality is a fundamental property of the particle that is frame-independent. A fermion can either be in a left-handed or right-handed chiral state. If a fermion is travelling in space, it maintains its chirality in all interactions except for interactions with the Higgs which can cause its chirality to flip via the Higgs mass mechanism. The charged-current weak force couples either an up-type quark with a down-type quark, or a lepton and its corresponding neutrino. One can therefore

arrange the left-handed fermions into a set of doublets, such as $\begin{pmatrix} q_u \\ q_d \end{pmatrix}_L$ or $\begin{pmatrix} \nu_l \\ l \end{pmatrix}_L$,

where the W^\pm bosons will couple between the upper and lower portions of the doublet. We can assign a value that indexes whether or not the particle is at the upper or lower portion of the doublet, which is referred to as the weak iso-spin $I_{3,L}$,

which will be $\begin{pmatrix} \frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}_L$. As right-handed particles are not included in this picture,

they are considered singlets assigned $I_{3,R} = 0$. As the label of which is the $I_{3,L} = \frac{1}{2}$ and which is the $I_{3,L} = -\frac{1}{2}$ portion of the doublet is completely arbitrary, the physics should not change upon the exchange of this label. This symmetry is referred to mathematically as the weak interaction belonging to the $SU(2)_{\text{weak},L}$ symmetry group,

which also represents the conservation of weak isospin through Noether's theorem [11]. The quark weak eigenstates are not the same as the quark mass eigenstates, which allows the weak force to interact between quark generations [12, 13]. Similarly, the observation of neutrino oscillations demonstrate that, similarly to the quarks, the weak eigenstates not being the same as the mass eigenstates, but rather a superposition of them [14].

The generators of the $SU(2)_{\text{weak},L}$ symmetry, as well as observation, requires the weak force be mediated by the charged W_1 and W_2 bosons, which mix to become the W^\pm bosons, and the neutral W_3 boson, all of which only interact with left-handed particles. There is separately another neutral current interaction boson B that mediates interactions with both left-handed and right-handed particles, and conserves a quantity known as weak hypercharge $Y = 2(Q - I_3)$ through the $U(1)_{\text{EW},Y}$ symmetry group. However due to electroweak symmetry breaking, The B and W_3 are mixed into states γ and Z , parameterized by a mixing angle θ , where Z now interacts with both left-handed and right-handed particles. The photon γ now no longer obeys $U(1)_{\text{EW},Y}$ symmetry but rather $U(1)_{\text{EM},Q}$ symmetry, where charge Q is conserved, and is what is considered the electromagnetic force.

The mathematical formulation of the SM is written out in the form of a Lagrangian. One cannot write terms that give particles masses in the classical sense as these terms automatically mix left-handed and right-handed components, which means it would not be possible for them to describe the behaviour observed where they interact differently with the weak force. However, interactions of the fermion with another field allow this separation. The Higgs mechanism circumvents this issue by replacing the particle mass in these terms with an interaction with another field, the result of which happens to be set at the mass of the particle. This other field is the Higgs field. The Higgs potential has a minimum, referred to as the vacuum expectation value (VEV), at a non-zero value such that the effective mass it imparts on fermions is non-zero. The different fermion masses are explained by different strengths of coupling with the Higgs field, known as Yukawa couplings. The Higgs boson is the particle that mediated interactions with the Higgs field.

As right-handed neutrinos (and left-handed antineutrinos) have not yet been observed in nature (although their existence has not been excluded), they have not been included in the SM. Note however that as the Higgs mechanism flips particle chirality, they would need to exist if neutrinos received mass via the Higgs mechanism.

Neutrino oscillation experiments have shown that do in fact have some mass, albeit small, however the mechanism through which they obtain mass is not yet known [6].

1.2 The Motivation Behind New Physics

The SM thus far describes observations in experiments remarkably well in many precision measurements, particularly in the QED sector such as in the electron magnetic moment [15]. However, there are still some open questions in physics that the SM cannot yet fully explain. Of these, two are particularly relevant to this work: the Higgs mass problem (also known as the Hierarchy Problem or the fine-tuning problem) and dark matter (DM) [16].

1.2.1 The Higgs Mass Problem

In the SM, the calculations related to any sort of process usually involve needing to calculate a cross-section or lifetime, which requires the calculation of a process matrix element. The process matrix element will be a sum over all the matrix elements of the Feynman diagrams that can contribute to the process. For example, if one considers the process of electron-positron scattering, there will be the different possible tree-level diagrams like the s and t -channels, and then also higher order diagrams that will involve closed loops. Note that the higher order diagrams typically contribute less to the final process matrix element as, since they will have multiple vertices. Each vertex carries a factor of the force "charge", which is typically a small number, meaning the final magnitude of their contribution will usually be small [7].

If one considers a single individual particle flying through vacuum, there will be loop corrections of every order, and for every particle that it can interact with, which depending on the type of particle can additionally include itself. These loops will occur infinitely many times over time, and so when one observes the particle, one does not observe the bare particle moving through space but rather this entire system. Instead of having to deal with this in terms of summing over infinite diagram matrix elements, one can instead reframe this as a correction to the mass of the simplest regular single particle diagram, where this would be the mass we would observe. This is referred to as mass loop corrections [17]. A schematic describing these loop corrections for the Higgs is shown in Figure 1.2.

Figure 1.2: The system of infinite single fermion loops occurring sequentially can be interpreted as a correction to the particle mass.

Of course it is not sensible to integrate over the infinite energy of the fermion in the loop. Once the fermion reaches more than the Planck mass, or equivalent energy of about 1.22×10^{19} GeV, the quantum field theory framework used in this calculation is no longer a valid model of reality as gravity becomes very important at this scale [19]. Evidently, this means that there must be drastically different new physics before the $\Lambda = 1 \times 10^{19}$ GeV scale, so it does not make sense to be integrating over energies larger than this. Instead, cutting off the integration at this value in a procedure known as regularisation [17], the corrections become: $\Delta m_H^2 \propto \int_0^\Lambda \frac{|k|^3}{|k|^2 - m_f^2} |k| \approx \Lambda^2 + m_f^2 \ln \left(\frac{\Lambda^2}{m_f^2} \right)$.

This essentially results in $m_{H,\text{Observed}}^2 = m_H^2 + \Delta m_H^2$ where $m_{H,\text{Observed}} = 125 \text{ GeV}$ but $\Delta m_H^2 \rightarrow (1 \times 10^{19} \text{ GeV})^2$, meaning that the bare original m_H must also be very large in magnitude, and also happen to cancel out the corrections almost perfectly to leave behind 125 GeV. This seems like a remarkable coincidence, about the equivalent of correctly guessing the exact number of insects living on Earth and only being off by about 125 [20]. This naturally raises suspicion about the "naturalness" of this cancellation, additionally because this divergence will occur with any fermion the Higgs couples to, whether discovered yet or not. This is known as a fine-tuning problem, and hints that there may be other particles out there that are interacting with the Higgs to help cancel out its diverging corrections in a more sensible way

(particularly bosons, which can have different loop diagrams which result in an opposite sign divergence) [21][17]. If this is the case, the Higgs itself can be used as a probe to try to find these new particles Beyond the Standard Model (BSM) as it couples to them.

1.2.2 Dark Matter

The dark matter problem arose when astrophysicists noticed that the velocity at which many celestial objects rotated around their galactic centres was faster than expected given the gravitational attraction one would expect from an estimation of their visible mass. Note that with telescopes, one can only observe luminous or absorbing matter. This discrepancy in what are known as galactic rotation curves strongly hinted that there must be some extra gravitationally-interacting *matter* that are in these galaxies that are non-luminous and non-absorbing, i.e. *dark* [22].

This is further strengthened from data deriving from gravitational lensing, where one can determine the amount of mass at a location by how its gravitational field affects light originating from an object behind it. Using this method on galactic collisions, particularly one now known as the "Bullet Cluster", astrophysicists had observed that the locations where the majority of the mass was observed, did not coincide to locations of where luminous matter was observed. Whereas the shape of the luminous matter distribution implies that the luminous matter from each of the colliding galaxies have interacted with each other in the collision, the "invisible" matter distribution implied that the "invisible" matter from each galaxy simply passed through each other and the luminous matter without interacting. The Bullet Cluster is considered the strongest evidence of electromagnetically non-interacting gravitationally-interacting dark matter. This additionally ruled out the possibility that the results of the galaxy rotation curves could have just been caused by gravity behaving differently at galactic scales [23].

However, despite knowing its density from measurements of the Cosmic Microwave Background (CMB) [24], not much is actually known about what dark matter actually is in terms of composition. One of the several proposed solutions to this problem is that dark matter is some sort of massive BSM fundamental particle. However, there are a number of conditions on these particles that can already be imposed, namely that they do not interact particularly strongly with the electromagnetic field

(i.e. neutral), they have mass, and they must be stable (otherwise we would observe evidence of their decays) [21].

In general this is not satisfied by particles in the SM. Neutrinos, although seemingly a good candidate on account of being very weakly interacting, stable, and neutral, are too light. Unfortunately this means that 1) they tend to be produced travelling at close the speed of light, meaning that they would travel too quickly to form the galactic structures observed, and 2) there are not enough of them to match the density of dark matter observed. This therefore means that whatever dark matter is, it has yet to be discovered [25].

1.3 The Motivation for Long Lived Particles

As of this writing, no widely accepted evidence for any BSM particles has been found. However, the vast majority of searches to date for BSM particles, at least at the LHC, work with the assumption that the BSM particle is short-lived (namely $\tau < 1 \times 10^{-15}$ s) or stable. This assumption is critical as it drives the reconstruction algorithms and even the detector design, making the majority of searches largely blind to discovering BSM particles if they happen to be long-lived [16]. There is not a precise definition of what is considered long-lived, as this depends on the detector sensitivity. A long-lived particle (LLP) can be effectively any particle that can be distinguished by the experiment from a particle that has no lifetime. ATLAS long-lived particle summary plots, when including signatures in all subdetectors, typically show long-lived particle lifetimes as being in approximately in the $1 \times 10^{-4} \text{ ns} < \tau < 1 \times 10^3 \text{ ns}$ range [26].

The majority of current experiments working on discovering LLPs involve the Large Hadron Collider (LHC). These can either be through the major experiments such as ATLAS, CMS, or LHCb, or through auxiliary experiments such as FASER or CODEX-b designed to increase the lifetime sensitivity of one of the major experiments [27]. There are however some other experiments also searching for LLPs that do not involve the LHC, for example in DUNE where one can look at non-standard decays to have sensitivity to much longer lifetime LLPs that may have been produced at the neutrino production facility [28].

However, there is not necessarily any particular strong reason for a potential BSM particle to be short-lived. Many SM particles are long-lived, such as muons, kaons or b

mesons. Furthermore, the fact that no new physics has been discovered yet by only looking for short-lived particles should in itself be a hint that perhaps we should be questioning that assumption and expanding our search to include that possibility [16].

The theoretical basis of the emergence of a long-lifetime particle lie in the theory parameters satisfying certain conditions. The equation for a particle lifetime is as follows, where the particle lifetime is equal to the reciprocal of the sum of the decay widths for every decay mode. Therefore, the more decay modes a particle has, the shorter lived it will be and vice versa.

$$\tau = \frac{1}{\sum_{\text{decay modes}} \Gamma_{\text{decay mode}}} \quad (1.1)$$

The formula for the decay width $\Gamma_{\text{decay mode}}$ is from Fermi's golden rule and includes a contribution from the density of the final states (i.e. available phase space of the decay products) as well as the matrix element \mathcal{M} of the interaction. For a decay into two particles, it will be proportional to the following, where X represents the original particle and a and b represent its decay products in an $X \rightarrow a + b$ decay. $|\vec{p}^*|$ represents the magnitude of the centre of mass frame momentum of either decay product.

$$\Gamma_{\text{decay mode}} \propto \frac{|\vec{p}^*|}{m_X} |\mathcal{M}|^2 \quad (1.2)$$

$$\propto \frac{\sqrt{(m_X^2 - m_a^2 - m_b^2)^2 - 4m_a^2 m_b^2}}{2m_X^2} |\mathcal{M}|^2 \quad (1.3)$$

$$(1.4)$$

From the $(m_X^2 - m_a^2 - m_b^2)^2$ term in the phase space contribution, it is evident to see that, given the constraint $m_X \geq m_a + m_b$, if the difference in mass between the original particle and its decay products is small, this term will become smaller. This means that there is less phase space for the decay (i.e. fewer "ways" in which the decay can occur). This reduces the decay width, makes the decay less likely, and prolongs the lifetime.

The matrix element term will usually be proportional to some model-defined vertex factor coupling between the particle and its decay product, which we can call g . Since

$\mathcal{M} \propto g$, a weaker coupling between the particle and its decay product makes its decay less likely, reducing the decay width and lengthening the lifetime.

Additionally, if one of the decay products was instead some mediator, for example $X \rightarrow a + b$, but where $b \rightarrow c + d$, then the propagator for b enters the equation for the matrix element, so $\mathcal{M} \propto g \frac{1}{q_b^2 - m_b^2}$. Therefore, if the mediator b is forced to have a virtual mass q_b very different than its rest mass m_b , known as being "offshell", this low value of the propagator reduces the probability of decay and lengthens the lifetime.

Therefore, the four factors that can result in any particle from any theory being long-lived are as follows:

- The particle has few decay modes;
- The particle has small-mass splittings with respect to its decay products, and therefore less phase space for the decay;
- The particle has small couplings between it and its decay products;
- The particle decays via a very off-shell mediator.

The difficulty in searching for BSM particles may hint that their coupling with the SM is weak, which would naturally lead them to have longer lifetimes, further motivating long-lived particles as a search space [16].

These properties that lead to particles being long-lived are not specific to any model, but can arise in any model if it just so happens that the model parameters result in the above phenomena. This means that the advantage of long-lived particle searches are that, as the signatures are relatively low background and fairly general, searches can be model-agnostic, allowing them to be interpreted in a wide variety of models. This additionally has the advantage and not needing to specify any particular assumptions in terms of the model or even the mechanism for which the particle becomes long-lived [16].

However some examples of models still need to be chosen for the purpose of analysis design, even if the search is meant to be general and not to target only these models specifically. These are known as benchmark models. Examples of typical models used as benchmarks within long-lived particle analyses for the purpose of analysis design relevant to this work are described in the following sections.

1.4 Selected Benchmark Models

There are a plethora of proposed models to solve these issues, including supersymmetry (SUSY) and the Higgs/Scalar Portal model, all of which propose additional BSM particles that would interact with the Higgs in such a way to help cancel out some of the corrections. Some of these particles may additionally be candidates for dark matter [16].

1.4.1 Higgs/Scalar Portal

The Higgs/Scalar Portal model is not a specific defined model but rather an effective simplified phenomenological scenario. The primary assumption made in this scenario is that dark matter exists (possibly in some dark sector with other particles that may or may not be considered) and obtains its mass from the Higgs mechanism, therefore meaning it can interact with the SM at least via the Higgs [29]. It is typically assumed that the new dark sector particles are bosonic, such that they can help cancel out the SM fermionic loop corrections in order to solve the Higgs mass problem. The Higgs Portal model is sometimes called the Scalar Portal model, as the Higgs that connects the SM with the dark sector may not necessarily be a SM Higgs, but could also be some other as-of-yet undiscovered scalar that could be a type of Higgs [29]. However for the purposes of this work, we will only consider a SM Higgs Portal.

This means that if the dark sector particle being searched for has less than half the Higgs mass, any normal SM Higgs boson (which is able to decay into a particle-antiparticle pair of any particle it couples with) produced at the LHC should be able to decay into a particle-antiparticle pair of the dark sector particle. In theory one could also consider dark sector particles with more than half of the Higgs mass, however the production cross-section for these is likely to be very low due to the small Higgs width [10], as well as the width of the dark sector particles also likely be small if they are long-lived. Precision studies have shown that the upper bound on the branching ratio of Higgs to undetected decays is up to 0.15 at 95% confidence level meaning that there is plenty of room for discovery in this scenario [30], especially considering that the Higgs itself was in part discovered via the $H \rightarrow \gamma\gamma$ process with a branching ratio of 2.270×10^{-3} [31][4].

This is useful in terms of searches as one can trigger on the Higgs production process without needing to assume as many details about the specific particle or its decay process, other than its mass and its coupling with the Higgs. This minimal and well-motivated scenario makes it incredibly useful as a benchmark model that can be used across analyses [26]. The Higgs Portal model is a good step forward in decoupling the signature from a specific model, especially since the assumptions made by the Higgs Portal model are those that are already included in many of the widely-used BSM models such as supersymmetry [17].

This is particularly helpful for long-lived particle searches. Since long-lived particles can come from any model whose parameters happen to satisfy the conditions outlines in Section 1.3, there are many different theories and simplified supersymmetry models that are used by various analyses, some of which can be very simplified and limited to a very specific area of parameter space. This results in a difficulty in directly comparing results of different analyses. However if an analysis can be interpreted via the Higgs Portal model, this allows for a very clear benchmark and comparison across a wide variety of signatures in different analyses [26].

Dark matter is generally accepted to be stable [25], however many theories theorize an entire dark sector, where there may be multiple higher generation dark sector particles, of which the lightest stable particle would be dark matter [21]. Long-lived particles being searched for at the LHC would therefore need to be one such unstable higher generation dark matter particle in the dark sector. The motivation for searching for such a particle are that the signature would be very clear in the detector (as opposed to dark matter itself which as it is neutral would only be detectable as a missing transverse energy), and that the matter particles in the SM themselves come in generations [29]. In theory one should also place constraints on the mass of the dark matter in the Higgs portal model and its coupling with the Higgs in order to ensure that it is compatible with the relic density. However these constraints are not particularly restrictive, hence why they are not usually considered in depth in analyses that target these models[32].

As for the Higgs production processes themselves, there are 5 different primary Higgs production processes at the LHC. The top three in production cross-section for $\sqrt{s} = 13$ TeV are gluon-gluon fusion at 43.92 pb, vector-boson fusion (VBF) at 3.727 pb, and the WH processes at 1.362 pb. The other modes are ZH and ttH production [33]. The Feynman diagram for these processes are shown in Figure 1.3. Most of these processes additionally produce objects alongside the Higgs that can be easily triggered

on in the detector. A fraction of gluon-gluon fusion events can produce initial state radiation jets that can be triggered on, VBF will always have jets that can be triggered on, and WH can produce jets or missing transverse energy (MET) and a lepton that can all be triggered on, depending on the W decay mode.




Figure 1.3: Feynman diagrams for the various Higgs production process. The two partons on the left-side of the diagram originate from the two different protons being collided.

There may be extensions to this most simplified scenario with additional assumptions. This work will focus on a particular such scenario where the dark sector particle being searched for will be a higher generation of dark matter that happens to be a long-lived scalar S that decays back into some SM products. It is assumed to be a scalar as these would help cancel out loop corrections in the Higgs mass problem, and that it is long-lived as its coupling to the SM particles it decays into eventually is weak. The scalar is also assumed to have Yukawa couplings with SM fermions, much like the scalar Higgs boson, and will therefore preferably decay into heavier particles such as heavy quarks. Another assumption is that the portal itself is a SM Higgs [26].

1.4.2 Supersymmetry (SUSY)

Unlike the Higgs/Scalar Portal scenario, supersymmetry (SUSY) is a model framework that predicts the existence of a whole suite of very specific particles in its "dark sector". A primary motivation of the theory lies in attempting to solve the hierarchy problem. With the Higgs mass corrections, bosonic interactions with the Higgs diverge with the opposite sign as fermions do [17].

Therefore, if every fermion had a bosonic counterpart with the same mass and vice versa, then these loop corrections would cancel out perfectly. These new supersymmetric partners to SM particles are denoted by placing a tilde on top of the symbol, and adding an "s-" prefix for sfermions (superpartners of fermions) and an "-ino" suffix for bosinos (superpartners of bosons) to their names. Each quark q , lepton l , and neutrino ν would have a squark \tilde{q} , slepton \tilde{l} , and sneutrino $\tilde{\nu}$ scalar counterpart and the gluon g has a spin- $\frac{1}{2}$ gluino \tilde{g} counterpart. Supersymmetry occurs before electroweak supersymmetry breaking, hence superpartners for the electroweak bosons are the Winos and Binos for W^1, W^2, W^3, B rather than for W^\pm, Z, γ . It turns out that in order for supersymmetry to properly model the required symmetries (i.e. avoid what is known as a gauge anomaly), additional Higgs are required. In the "SM" sector (as opposed to the supersymmetric sector), the Higgs sector would need to comprise the regular SM spin-0 Higgs, two charged Higgs, and an additional neutral Higgs. These would all have their corresponding superpartners. The two positive, two negative, and four neutral higgsinos and electroweakinos later mix, resulting in observable states of two positive and two negative charginos $\tilde{\chi}_1^\pm, \tilde{\chi}_2^\pm$ and four neutralinos $\tilde{\chi}_1^0, \tilde{\chi}_2^0, \tilde{\chi}_3^0, \tilde{\chi}_4^0$ [17].

Proposing such "superpartners" to every SM particle results in a new symmetry: that if all fermions were exchanged with their bosonic counterparts and vice versa, then physics remains unchanged. This is the symmetry of supersymmetry [17].

Of course this supersymmetry is broken via some unknown mechanism, otherwise one should have already observed these superpartners at the exact same mass as the SM particles. This symmetry breaking presumably pushes the masses of the superpartners higher to where we have not been able to search yet. Although this removes the perfect cancellation of the divergent loop corrections, the prediction is that as long as the mass difference between the SM and its superpartner is not too large, there is still adequate enough cancellation for the loop corrections to remain "natural" [17].

There are multiple possible SUSY theories, some with additional particles to the ones listed. However, usually the Minimal Supersymmetric Standard Model (MSSM), which would be the minimal SUSY extension to the SM, is considered in searches. There is still a very large number of free parameters (about 105) in the model that would need to be tuned and fit to experiment, particularly particle masses, couplings, and mixings [17]. It is not feasible to create a scan over the full parameter space where one varies each of these parameters in order to compare to data. Typically in searches, simplified models are used where only the Feynman diagram in question are considered. It is assumed that all other particles are at either very high mass or very low couplings such that they are effectively decoupled from the process in question and that the branching ratios are 100%.

The majority of SUSY searches target a class of SUSY theories that contain the assumption that R-parity is conserved. This is known as R-parity conserving SUSY. R-parity is defined as

$$R \equiv (-1)^{3B+L+2S} \quad (1.5)$$

where B is the baryon number, L is the lepton number, and S is the spin of the particle. This definition results in SM particles having $R = +1$ and their superpartners having $R = -1$. R-parity conservation therefore corresponds to the "SUSY/SM particle number conservation", where the number of particles minus antiparticles of the SM and the number of superpartners minus antisuperpartners remains conserved across the diagram. R-hadrons are consequently hadronic bound states that contain both supersymmetric particles and regular quarks [34].

Imposing this conservation is attractive as this automatically guarantees that the proton would remain stable, which is what is observed. Additionally, R-parity conservation would mean that the lightest stable particle (LSP) of the model (typically the $\tilde{\chi}_1^0$) would also be stable, making it a dark matter candidate [17].

However, R-parity is not a necessary condition for proton stability as couplings can still be designed such that proton stability remains but certain R-parity violating (RPV) decays can still occur [17]. Due to this, searches for R-parity violating SUSY models are still conducted at LHC experiments such as ATLAS, and can come with unique signatures such as large numbers of leptons or jets in the final state [34]. In these RPV cases, there may not necessarily be a dark matter candidate.

Chapter 2

Experimental Background: The LHC and the ATLAS Detector

“Well, [LHCb’s] flavour anomalies are gone, [ATLAS and CMS’s] SUSY is dead, who knows what ALICE does...”

— Jenny Smallwood, PhD Particle Physics, University of Oxford

2.1 The LHC at CERN

One approach to searching for new physics involves trying to produce new undiscovered particles through high energy collisions of known and easier to acquire particles by accelerating them and then colliding them into each other [35]. The world’s highest energy particle accelerator is the Large Hadron Collider (LHC) at CERN in Geneva, Switzerland, which for the most part, and more importantly for the purposes of the analyses in this work, collides protons. The LHC does so by accelerating protons in bunches of approximately 1.2×10^{11} to energies of 6.5 TeV (during the Run 2 data taking period from 2015 to 2018), in opposite directions in a 27 km underground accelerator ring, such that they collide at four different interaction points around the ring. The protons’ combined centre-of-mass energy during the collision is therefore $\sqrt{s} = 13$ TeV for Run 2 [36].

However, one cannot simply inject protons at rest directly into the LHC and have them accelerate immediately to 6.5 TeV. There is instead the pre-accelerator

complex which comprises a chain of historic and progressively increasing energy accelerators that protons undergo before entering the LHC, described in Figure 2.1. The protons begin in the form of hydrogen gas in a bottle. The protons are extracted from the gas through ionization in a duoplasmatron, before being injected into a linear accelerator (LINAC 2 during Run 2). LINAC 2 accelerates the protons to 50 MeV. The protons then enter the booster rings which accelerate the protons to 1.4 GeV, then the Proton Synchrotron (PS) which accelerates them to 26 GeV, before they enter the 5 km circumference Super Proton Synchrotron (SPS) which accelerates them to 450 GeV, before they are finally injected in opposite directions into the LHC [37].

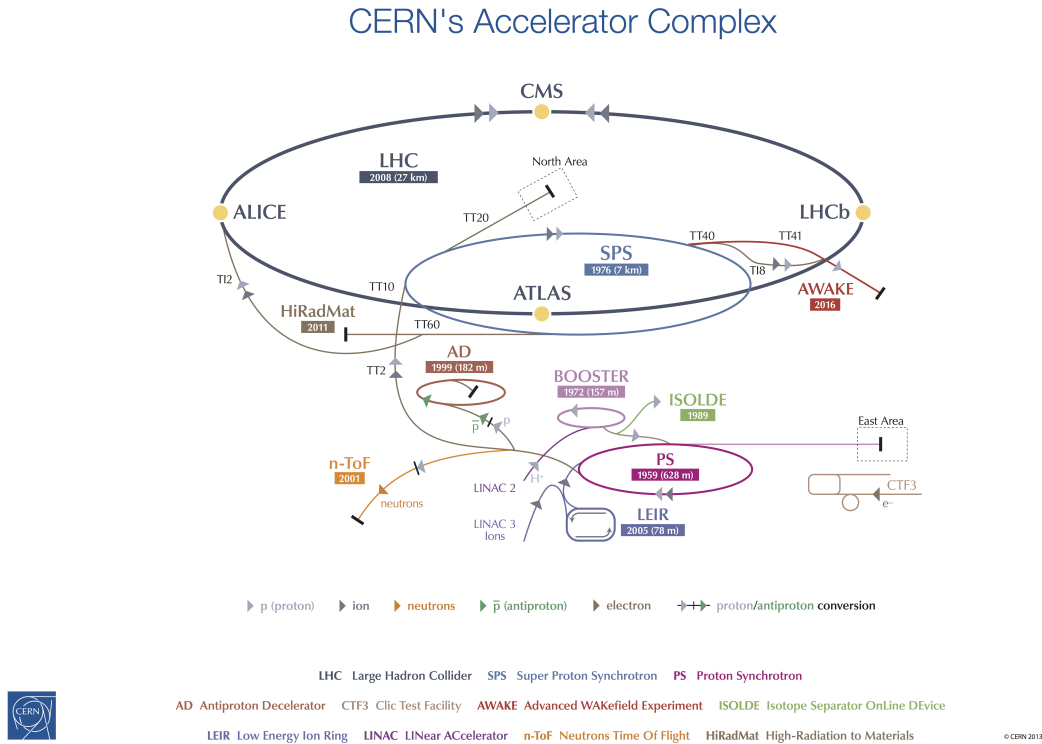


Figure 2.1: A diagram of the CERN accelerator complex. The protons start in a hydrogen bottle before being injected into LINAC 2, before finally ending up in the LHC [38].

Situated at the four interaction points are detectors, each being run by a different independent collaboration, and each being designed to be specialized for investigating different aspects of new physics. These experiments are ATLAS, LHCb, CMS, and ALICE. ATLAS and CMS are more general purpose detectors that mainly focus on searching for new BSM particles from various different theories such as SUSY. LHCb specializes in the measurement of charge-parity violation in hadrons containing b -

quarks, trying to discover evidence of new physics from finding anomalies in precise measurements. ALICE specializes in heavy ion and nuclear physics [35].

From 2015-2018 at the LHC, known as Run 2, bunch crossings occurred every 25 ns, with on average about 33.7 proton-proton interactions within each bunch crossing [39]. The data produced from these bunch crossings are called events. Since it is not possible to write out all of these events to memory, and the vast majority of these collisions do not produce anything particularly interesting in the context of the physics program, each experiment decides which events to record by the use of a trigger. These triggers usually search for some general interesting activity in some part of the detector, such as high missing transverse energy (MET) or high energy jets or muons [40]. This means that in an analysis, typically one must start by defining a theoretical process to target in the search, then figure out which triggers would be triggered by that process. If a suitable trigger is not available, one would need to develop a new trigger, however this would only be possible to use during future data-taking. One would then typically look at the number of events that fired the trigger, select a signal region, and see if one observes more events or a different shape than expected from the SM or noise processes in that region.

Because the number of events that will be seen for each process is different, the amount of data recorded by the LHC is stated in terms of integrated luminosity, $\int \mathcal{L} dt$, where \mathcal{L} is the instantaneous luminosity at a given time. The luminosity encapsulates information about the particle beams such as beam alignment, remaining protons in the bunch, beam width, and other properties [41].

If for example one was to repeatedly throw balls to land within a circle on a wall, the cross-section σ of the process would be the area of the circle. However the number of balls N that actually land in the circle depends on how well the balls are thrown, and this can change over time. The luminosity encapsulates this information on how well the ball is being thrown, such that if the balls were being thrown similarly at different circles of different cross-sections, the number of hits at each circle can be predicted.

The instantaneous luminosity \mathcal{L} can be calculated via a test process with a cross-section σ where the number of events N is counted over a small time period, as $\mathcal{L} = \frac{1}{\sigma} \frac{dN}{dt}$. This way, the number of events N that will be observed in the data for any

given process with a cross-section σ will be given by

$$N = \sigma \int \mathcal{L} dt. \quad (2.1)$$

The integrated luminosity of the data ATLAS recorded that was good for physics for all of Run 2 in 2015-2018 was 139 fb^{-1} [39]. Run 1, which ran from when the LHC started in 2010 until 2012, delivered a luminosity of approximately 25 fb^{-1} at around $\sqrt{s} = 7\text{-}8 \text{ TeV}$ [41]. Run 3 began on July 2022 and is as of this writing running with $\sqrt{s} = 13.6 \text{ TeV}$ [42]. Unfortunately, due to the relatively longer completion time-scales of LLP analyses, no Run 3 data is able to be included in this work and the analyses discussed focus only on Run 2 data.

2.2 The ATLAS Experiment

ATLAS (A Toroidal LHC ApparatuS), other than being a rather unconvincing acronym, is a cylindrical-shaped 44 m long, 25 m in diameter [43], 100 m underground [36], 7000 ton general purpose detector surrounding the beam-pipe at one of the interaction points on the LHC [43].

For reference, the coordinate system used in ATLAS has the z -axis parallel to and in the centre of the beam-pipe. The y -axis is oriented vertically where the positive direction is towards the sky. The x -axis is horizontal and perpendicular to the beam-pipe where the positive direction is towards the centre of the LHC [43].

The azimuthal angle around the z -axis in the x - y -plane is ϕ , where $\phi = 0$ points horizontally to the centre of the LHC, and $\phi = \pi/2$ points to the sky. The polar angle θ from the x - y -plane towards the z -axis is usually written in terms of the pseudo-rapidity

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right]. \quad (2.2)$$

Note that $\eta = 0$ refers to an angle perpendicular to the beam-pipe, whereas the angle approaches the beam-pipe as $\eta \rightarrow \infty$ [43]. Note that very often, quantities are expressed in terms of their transverse component, which is the component of their vectors in the x - y -plane.

The design of the ATLAS detector is such that the proton-proton collisions occur in the centre of the detector so that the resulting daughter particles emanating from

the primary vertices (PVs) can be detected by the detector, which comprises multiple specialized subdetectors arranged in layers [43].

The inner-most subdetector of ATLAS is the *Inner Detector* (ID). The purpose of this is for the precise determination of the trajectories of charged particles such that the charge and momentum of the particles can be measured. Note that neutral particles cannot be detected here. Next comes the *electromagnetic* (EM), then the *hadronic calorimeters* which stop and determine the energy of hadronic particles. The EM calorimeter also stops and determines the energy of electrons and photons. Finally the outermost layer is the *muon spectrometer* (MS), which usually only muons can reach, which is meant to track and trigger on muons, obtaining a charge and momentum measurement for them. Note that neutrinos and any non-interacting BSM particles would fly through the entire detector without interacting with anything. A diagram of the ATLAS detector is shown in Figure 2.2 [43].




Figure 2.2: A diagram of the ATLAS detector [43].

In order for the charge and momenta of particles going through the detector to be measured, the portions of the detector that measure trajectories are immersed in magnetic fields. The magnetic fields of strength B bend the trajectories of the particles. Their direction of bending will reveal charge information, and the radius of bending r in the plane perpendicular to the field will reveal information about the momentum component perpendicular to the B field as per the Lorentz formula $p_{B\perp} = |qrB|$, where q is the charge of the particle. The ID is immersed in a 2 T solenoidal magnet situated

between the tracker and the calorimeters. This magnetic field is along the beamline, causing the particles to curve in the circular plane of ATLAS, measuring the p_T of those particles. The MS has a series of eight toroidal coils going through it in the barrel section (hence the "toroidal" in the name of ATLAS) with a smaller version of those eight coils on either end of the barrel in the end-caps, providing a field up to 3.5 T. This magnetic field goes around ATLAS in a circle that goes through all of the toroids. This means that the muons would actually get bent in a direction orthogonal to how they were bent in the ID, providing information about the muon p_z , allowing for very accurate muon momentum measurements [43].

Readout of individual sensor elements of the detector will usually be processed into packets by readout chips on the sensors or a group of sensors themselves. These then transmit the data out of the detector via optical links. The L1 hardware trigger is then done underground, while the HLT software trigger occurs on farms on the surface. Raw data from events that pass the triggers are transmitted to the data centre for storage, and reconstruction then can then occur on the worldwide computational grid [43].

The following subsections will outline the different sub-detectors of ATLAS. They are all detailed in the ATLAS detector technical specification document [43].

2.2.1 The Inner Detector (ID)

The inner detector (ID) is itself composed of three different sub-detectors, similarly arranged in concentric layers surrounding the interaction point, extending 5.3 m in length and 2.5 m in diameter. The innermost layer are the high-resolution silicon pixel detectors (PIX), followed by the silicon microstrips detectors (SCT), then finally the coarser transition radiation tracker (TRT) [43].

The barrel portion of pixel detector, critical to DV reconstruction, comprises four layers starting at $r = 33.25$ mm, 50.5 mm, 88.5 mm, and 122.5 mm. The innermost layer is called the Insertable B-Layer (IBL) and was installed for Run 2. The other pixel layers are, from inside out, the B-Layer (BL), Layer 1 (L1) and Layer 2 (L2). The IBL runs from $z = -331.5$ mm to 331.5 mm while the other layers run from $z = -400.5$ mm to 400.5 mm. The sizes of the pixels are 50 μm in the $r\phi$ direction and 400 μm in the z direction for the outer layers, and 50 μm by 250 μm for the IBL. In addition to these

barrel detectors there are three pixel endcap detectors from $r = 88.8$ mm to 149.6 mm at $z = 495$ mm, 580 mm, and 650 mm [44] [45] [43].

There are four SCT barrel layers from $z = -749$ mm to 749 mm at $r = 299$ mm, 371 mm, 443 mm, and 514 mm, and multiple SCT end-cap layers up until $z = \pm 2720.2$ mm. These, along with the much larger TRT, allows for the detection of particles over 0.5 GeV in p_T within the pseudorapidity range of $|\eta| < 2.5$. There are two layers of strips per stave, with the strips being offset by 40 mrad in order to increase resolution in the $r - \phi$ direction [43].

The silicon-type detectors work by just being a pn-junction. A pn-junction is simply a junction between p-doped silicon and n-doped silicon. p-doped silicon has atoms implanted into the crystal structure that have one less electron than silicon does in their valence shell. n-doped silicon is the same except with an extra electron. When they are placed together a depletion area appears near the interface where electrons migrate from the n-doped side to fill the holes in the valence shells in the p-doped side, leaving positive ions in the n-doped side and negative ions on the p-doped side. The pn-junction for the current ATLAS detector is arranged in the sensor such that the bulk of the sensor is n-type and there is an implant of p-type silicon on the top surface. A reverse bias voltage is put across this pn-junction to expand the depletion region to cover as much of the sensor depth as possible. Note that after sufficient radiation damage, the n-type silicon in the bulk will turn into p-type silicon, requiring the voltage be inverted as the depletion area will now be between the newly p-type bulk and the n^+ -type (i.e. doped with enough n-dopant to the the point of being conductive) implant on the backplane of the sensor. The depletion area, due to the ions, has an electric field across it. This means that when charged particles pass through the depleted region of the bulk of the material creating electron-hole pairs are generated, these electrons and holes will be separated by the field and migrate to their respective sides of the wafers where they can be read out. However, of course the contacts and implants on the top of the wafer are separated into separate individual pixel or strip metal contacts to allow for only individual elements to have a signal when a particle is passing though the region below it. Note that there will be a small leakage current that still passes through the depletion region during normal operation, and electrons and holes will also be produced via thermal noise. This requires an insulating layer placed between the silicon and the readout metal to stop current, as well as requires the entire PIX and SCT to be cooled to between -5°C to -10°C in order to reduce thermal noise [43].

The TRT comprises many closely-packed straws filled with gas, each with a small wire going through the centre of it with a voltage difference between it and the straw. The space between the straws is filled with a material that changes rapidly spatially, such that particles traversing it emit transition radiation, which are photons emitted when relativistic particles traverse material boundaries. When a charged particle and its photons reach a straw, they produce electron-hole pairs in the gas, where the electrons are collected by the wire in the centre of the tube for readout. A diagram of the inner detector is shown in Figure 2.3. One of the major advantages of the TRT is the ability to distinguish between electrons and pions, as these particles have different amounts of transition radiation due to their different masses [43].




Figure 2.3: A cross-section diagram of the inner detector of ATLAS [46].

Note that since the p_T of the track is determining its radius of curvature via its arc height, its accuracy will depend on how accurate the arc height can be measured. High p_T tracks will appear very straight and with a small arc height, and as these are difficult to distinguish, these will be measured with less accuracy than low p_T tracks. In general, $\frac{\sigma_p}{p} \propto p$ [43].

Note that due to the planned upcoming high-luminosity LHC (HL-LHC) in 2028, the ID will need to be replaced to withstand the new harsher radiation environment. This new inner tracker (ITk) will comprise only silicon pixel and strips layers, made with

higher resolutions. Additionally they will be made with a p-type bulk and an n-type implant to avoid type inversion during radiation damage. [47]. Further details on these new strips sensors, related to the qualification task, are covered in the qualification task chapter of Chapter 5.

2.2.2 The Electromagnetic and Hadronic Calorimeters (ECAL and HCAL)

The calorimeters measure particle energy by placing dense material in its path to cause the particle (with the exception of muons who still survive the calorimeter if its energy is greater than 3 GeV) to lose all of its energy via hadronic and electromagnetic interactions. These hadronic and electromagnetic interactions create showers of secondary particles, which are then detected. The number of these secondary particles that is observed is what gives information about the energy of the original particle. As higher energy particles create more more secondary particles with, the uncertainty on the energy measurement is dominated by the Poissonian uncertainty of counting the secondary particles, hence $\frac{\sigma E}{E} \propto \frac{1}{\sqrt{E}}$ [43].

There are two technologies used in ATLAS for calorimetry. Tile calorimeters are used in the HCAL barrel, while liquid argon calorimeters (LAr) used for the ECAL, HCAL endcaps, and the forward calorimeter (another calorimeter that is placed in the forward direction) [43].

The tile calorimeter is made of several layers of steel and plastic scintillators. The steel converts the particle energy to secondary particles. The plastic scintillator contains compounds that absorb energy from passing secondary particles, and then fluoresce that energy into UV photons. These photons are led out of the scintillator in wavelength-shifting optical fibres that turn them into visible light. These pulses of visible light can then detected via sensitive photomultiplier tubes, which allows the number of secondary particles to be inferred [43].

The LAr technology has layers of metal that are separated by honeycomb spacers. These spaces are filled with liquid argon. An electrical current can be measured when secondary particles ionize the liquid argon. The layers are arranged in an accordion structure to ensure that there are no gaps through which original particles can travel without being stopped by the metal. The mean interaction length in a metal for a particle to interact electromagnetically will have a dependence on the number of

charged nucleons. However the mean hadronic interaction length only depends on the nucleon density. This means that different metals will have different electromagnetic and hadronic interaction lengths. One can therefore select a material with a shorter electromagnetic interaction length for the ECAL such as lead, and one with a shorter hadronic interaction length for the HCAL such as copper. Tungsten has both a lower electromagnetic and hadronic interaction length than the cheaper copper and lead, and so it is used for the forward region where there is a higher expected particle flux [43].

2.2.3 The Muon Spectrometer (MS)

The muon spectrometer is comprised of various different detector technologies. For the bulk of the barrel and endcap, the primary precision tracking detector are the monitored drift-tubes (MDTs). Layers of faster detector technologies are in most cases placed on either one or both sides of the MDTs in order to provide triggering capabilities. These are the resistive plate chambers (RPCs) in the barrel by, thin gap chambers (TGCs) in the big wheel of the outer end-cap, and the cathode strip chambers (CGCs) in the small wheel of the inner end-cap [43].

The muon drift tubes are a series of long 3 cm-diameter aluminium alloy tubes with a wire through its centre at a large voltage difference with respect to the tube. The tube is filled with and argon/carbon dioxide gas mixture, such that when a muon passes through the gas, it ionizes it and causes electrons to drift to the wire and be detected as the signal. The time it takes for the drift to reach the wire will depend on the radius of closest approach of the muon to the wire within the tube, allowing for an accurate tracking resolution. Unfortunately the drift time is relatively large (with a maximum drift time of 700 ns), meaning this cannot be used to trigger on events, and that there is some dead-time when a muon is detected [43].

It is for this reason that the multiple other faster detector technologies are used for triggering purposes, primarily by shrinking the distance that is needed for electrons produced from the muon ionizing the gas to reach the electrode under high voltage. RPCs have two resistive parallel plates, with the readout strips on the plates divided into strips, with a gas-filled 2 mm gap in between them. TGCs are a set of wires arranged perpendicular to strips in a gas chamber. The CSCs are similar to the TGCs

with chambers filled with gas, with perpendicular wires and strips at high voltage [43].

The small wheel of the inner end-cap is being replaced for Run 3 by the new small wheel (NSW), which instead of MDTs and CSCs, uses micromegas (MMs) and small-strip TGCs (sTGC). sTGCs are very similar to TGCs except most notably for a much smaller strip pitch. MMs comprise a planar cathode at -800 V and a readout electrode divided into strips at 0 V, separated by a gas gap of 5 mm. A mesh is placed 128 mm from the electrode, held at -550 V. The large area between the mesh and the cathode is the drift volume where electrons will relatively slowly drift to the mesh. Once the electrons hit the mesh the small gap and large voltage causes amplification via an avalanche of electrons, resulting in a fast readout [48].

2.3 Object Reconstruction

The detector will yield raw data largely in terms of detector elements that register hits in an event. This raw data needs to be reconstructed into physics objects that are useful for analysis, for example tracks, muons, electrons, or jets. This usually involves first forming seeds or clusters from the hits, and then using those to guide the search for other hits that may be part of the same object. Once All raw detector information involved with an object are found (for example hits or clusters), the reconstructed object can then be determined or fitted (for example tracks or jets). From this, information about these reconstructed objects, such as p_T or energy, can be calculated. These reconstructed objects then need to be calibrated to account for imperfect detector response. This section provides an overview of object reconstruction relevant to this work, namely jets, missing transverse energy (MET), and displaced vertices.

2.3.1 Jets

Quarks and gluons cannot exist for very long on their own and will undergo reactions to end up in colourless bound states within the timescale of about 1×10^{-24} s [49]. This usually results in the production of many hadronic, composite particles that will be travelling in the general direction of the original particle. In the ATLAS detector, this will appear as a spray of collimated hadronic particles, known as a jet. There are a

variety of reconstruction algorithms that can be used to define which particles belong to a jet, however more importantly, this means that if the final state one is looking for involves a quark or gluon, this will manifest itself in the detector, and therefore the data, as a jet [43].

Charged hadrons in the jet will leave tracks in the ID, and the hadrons themselves will be stopped and converted into energy deposits in the hadronic calorimeter which it measures. Jet reconstruction algorithms are used to cluster hits in the detector in order to identify the entire object as a jet. Two main algorithms are particularly important to this work: "EMTopo" and "track jets".

The calorimeter is divided into individual cells. The EMTopo algorithm begins with grouping these cells into topological-clusters. This starts with a seed cell with a large energy and adding adjacent cells above a certain threshold. This is then repeated until no other adjacent cells pass the threshold, after which all adjacent cells are added regardless of their energy. The anti- k_T algorithm [50] is then used to cluster together neighbouring topological-clusters [51]. Note that in analyses where a long-lived particle decays hadronically, one prefers such an algorithm that only relies on calorimeter deposits as there may not be tracks that are reconstructed in this case. This also allows for a useful long-lived particle signature: the trackless jet, where calorimeter deposits hint at the presence of a jet but there are insufficient corresponding tracks in the ID. Note however that it is assumed that the jets originate from the primary vertex, meaning that there may be mis-measurements if the jet may originate from a pileup vertex.

The same number of particles produced in the calorimeter when a particle hits can correspond to a particle of a certain energy interacting electromagnetically, or a particle of a different energy interacting hadronically. It is not known a priori which case is being measured. In ATLAS, it is assumed that all interactions are from electromagnetic processes. Additionally, as the detector is comprised of layers of showering material and detecting material, some of the secondary particles may be absorbed in the showering material and not be detected. These processes bias the energy measurement of the jet, hence calibrations are needed to recover accurate measurements [51].

In contrast, the track jet algorithm uses only track information. The anti- k_T algorithm [50] is run on standard tracks satisfying certain requirements. Each track is associated with only one proton-proton interaction vertex each before running the al-

gorithm, and the algorithm is run on all proton-proton interaction vertices, regardless if they are identified as pileup or not.

The anti- k_T algorithm tries to cluster objects to form jets that are within the size of a target cone parameter $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$. The algorithm begins by calculating the weighted distances between the entities being clustered (tracks in the case of track jets or cell clusters in the case of EMTopo jets) $d_{ij} = \min(p_{T,i}^{-2}, p_{T,j}^{-2}) \frac{(\Delta R)_{ij}}{(\Delta R)_{\text{target}}}$. Here i and j are the indices of the two entities being compared, p_T is the transverse momentum of the entity, $(\Delta R)_{ij}$ is the ΔR between the entities, and $(\Delta R)_{\text{target}}$ is the target ΔR of the final jets (i.e. the cone parameter). The distances between the entities and the beam are also calculated, as simply $d_{iB} = p_{T,i}^{-2}$. The entities are combined iteratively based on which pairs have the smallest d_{ij} . If d_{iB} is the smallest entity, then the cluster is called a jet and removed from the list. This is repeated until no entities are left. A distinguishing feature of the anti- k_T algorithm in particular are the negative exponents on the transverse momentum. This results in the prioritisation of higher transverse momentum entities being merged first. In this analysis, both the track jets and EMTopo jets use a cone parameter of 0.4 [50].

2.3.2 Missing Transverse Energy (MET)

The subdetectors of ATLAS rely on the electromagnetic and hadronic interactions to detect particles. This means that particles that do not interact electromagnetically or hadronically, for example neutrinos, will not leave any traces in the detector. However it is still possible to detect evidence of the presence of such "invisible" particles in an event due to momentum conservation. In the transverse direction, since the collision begins with no transverse momentum, momentum conservation dictates that the final state should also have no transverse momentum. ATLAS is a hermetic detector, in other words designed to surround and fully capture and measure all collision decay products (with the exception of neutrinos). One can therefore infer the amount of transverse momentum of the invisible particles in the decay products by measuring the amount of transverse momentum of the visible particles. Of course if there was more than one invisible particle in the event, one would only be able to detect their combined transverse momentum. Hence, the "missing transverse momentum" (also called "missing transverse energy" or MET) is an incredibly important quantity to determine in collision events, as invisible particles are important in many theories such as dark matter and this is the only way to infer their presence in the detector.

This is summarised in the following equation

$$p_{T,\text{miss}} = \left(\sum_{\text{invisible particles}} \vec{p} \right)_T = - \left(\sum_{\text{visible particles}} \vec{p} \right)_T. \quad (2.3)$$

Although momentum conservation holds in all directions, the reason why longitudinal momentum conservation is not considered is due to the initial momentum fraction of the protons carried by the actual partons that are interacting cannot be known a priori and would not necessarily be simply zero such as in the transverse direction. Note that one takes the transverse component of the vector sums, and not sum the scalar values of the transverse components [52].

The visible transverse momentum is calculated by taking the transverse component of the vector sum of all visible physics objects in the event. This includes muons which leave traces in all subdetectors, electrons detected in the ID and ECAL, photons detected in the ECAL, hadronic jets in the HCAL as well as some deposits in the ECAL and tracks in the ID, and finally the "soft" term which consists essentially of whatever remains in terms of tracks and calorimeter deposits not included in the previous objects. The objects are summed in this order, with all detector signatures associated with the object being removed after being added in order to avoid double counting. This is summarised in the following equation

$$p_{T,\text{miss}} = - \left(\sum_{\mu} \vec{p} + \sum_e \vec{p} + \sum_{\gamma} \vec{p} + \sum_{\text{jets}} \vec{p} + \sum_{\text{"soft"}} \vec{p} \right)_T \quad (2.4)$$

[52].

Note that only physics objects and tracks that can be associated with the primary vertex are counted in MET. This means that BSM LLPs can be reconstructed as MET in two ways: 1) the LLP if neutral and non-hadronically interacting lives long enough to fly out of the detector and is therefore essentially a normal invisible particle, or 2) the LLP decays late enough that the tracks of its decay products do not get reconstructed or associated with the primary vertex (i.e. treated like pileup) and don't enter the MET calculation.

2.3.3 Standard Tracking

The data recorded from the PIX and SCT are simply a record of which pixels or strips have been activated during an event. This means that the information available for reconstruction is that a charged particle of an unknown species passed through a specific volume element in the detector, known as a hit. However, such a reading could also be caused by noise. This hit data needs to be converted to more useful track information via reconstruction, where algorithms try to reconstruct the trajectories of the particles from the hit data [53].

Tracking begins by a seeding process. This involves searching for a series of "space points" from three different layers in either the PIX or SCT for a track seed in the PIX and SCT. These space points are a cluster of hits in a PIX layer, or a cluster of combinations of the two angularly offset strips in an SCT layer. A Kalman filter [Kalman] is used to determine all hits from the other layers are compatible with this track seed within the PIX and SCT within certain constraints. An ambiguity solving algorithm is then run to determine the best track to resolve cases such as, for example, when there may be more than one possible hit in a layer compatible with the track seed. Finally, TRT hits are considered and added to the track if compatible [53].

A second tracking pass, known as the "outside-in" pass as opposed to the previous "inside-out" approach, is then performed starting by seeding in the ECAL and reconstructing track segments in the TRT. These are then extended back to the PIX and SCT and associate any hits not part of existing tracks [53].

Tracks are parametrised by five parameters. There are the transverse and longitudinal impact parameters d_0 and z_0 , which is the track's point of closest radial and longitudinal approach to the beamspot. There is also the azimuthal ϕ and polar θ angles of the track calculated at the (d_0, z_0) point. Finally there is the charge over momentum magnitude ratio q/p [53].

A faster, modified version of standard tracking can be run by the HLT trigger over the entire ID or in a region of interest [54].

2.3.4 Large-Radius Tracking and Displaced Vertex Reconstruction

The ATLAS experiment in general assumes that electrons, positrons, protons, neutrons, muons, anti-muons, photons, and neutrinos are *detector-stable*, where their lifetimes

are long enough that they do not decay within the volume of the detector unless purposefully stopped by the detector [43]. More importantly, the ATLAS experiment generally assumes that all other particles produced by a collision are so short-lived that they essentially decay immediately at the primary vertex. This assumption is implemented by placing constraints on the tracks during reconstruction on the closest radial distance to the primary vertex $|d_0| < 10$ mm and the closest longitudinal distance to the primary vertex $|z_0| < 250$ mm, along with requirements on hit multiplicities. Hits that do not result in a track that satisfy these requirements are assumed to be noise and are usually ignored. As LLP decays may result in tracks that do not satisfy these requirements, some LLP signatures may use these ignored hits as a starting point for a secondary reconstruction step to target these scenarios, such as that used in this analysis [53].

There are many possible signatures in the ATLAS detector that LLPs can manifest as. Even in only the ID there are a multitude of possible signatures to search for, such as disappearing tracks [55], ionization loss [56], and displaced vertices. This thesis focuses on the displaced vertex signature. A displaced vertex is when a bunch of tracks seem to originate from a common point in the volume of the detector, and not from the primary vertex. This is due to the LLP travelling through the detector, which may or may not be visible to the detector, and then decaying into children tracks that are visible to the detector, such that the decay vertex is *displaced* from the primary vertex [57].

The d_0 , z_0 , and multiplicity requirements on tracks mean that the standard reconstruction algorithm used by ATLAS cannot reconstruct tracks that are from the decays of LLPs. In order to actually see the manifestations of these LLPs, a dedicated reconstruction algorithm needs to be run, called *large-radius tracking* (LRT). Following this, a secondary vertexing algorithm will need to be run to find the displaced vertices among the LRT tracks.

Since the standard tracking will result in some hits that it is unable to associate with a constrained track, LRT uses those hits to see if a less constrained track may be able to satisfy them. The LRT relaxes the constraints on the tracks to $|d_0| < 300$ mm and $|z_0| < 1500$ mm, and also relaxes the number of hits allowed to be shared by multiple tracks. LRT largely repeats the same steps as the standard tracking except with these relaxed constraints, and it does not do the outside-in pass. Note that these large-radius tracks will be treated as a separate single collection from the standard tracks in later steps [53].

For the determination of the displaced vertices from these tracks, another dedicated algorithm is used for secondary-vertex reconstruction, known as VrtSecInclusive (VSI) [58]. Only tracks with $p_T > 1 \text{ GeV}$ and $|d_0| > 2 \text{ mm}$ are considered, where the d_0 requirement helps ensure that we ignore tracks coming from the PVs. Note that this means that if the LLP itself is visible to the detector, it will not be included in the DV, only its children will be. Since there may be many fake tracks produced by the LRT due to spurious combinations of hits, several additional requirements are placed on tracks that may undergo the vertexing algorithm. These are defined in terms of the number of hits it sees in the PIX, SCT, and TRT portions of the detector. These requirements are even more stringent if the track has $p_T < 25 \text{ GeV}$ as there are more fake tracks at low momentum [57]. Tracks passing these strict requirements are referred to in this work as *core* tracks.

The actual vertexing process begins by first trying to find all possible 2-track combinations to determine "seed" vertices that satisfy a fit of χ^2 per degree of freedom < 5 . Further requirements are imposed on these seeds, namely that both tracks must not have hits closer to the beampipe than the vertex coordinate. The track momenta should also be such that the inner product of their sum with the vector pointing from the PV to the seed vertex coordinate should be positive, such that the vertex points away from the PV. Following this, these 2-track seed vertices are iteratively merged to form n -track vertices. Seed vertices are only merged if the distance significance between the two vertices (i.e. the distance divided by its uncertainty), $S^2 \leq 100$ [57].

However since the requirements for the core tracks that form the seeds are quite strict, the algorithm also allows for tracks to be *attached* to the vertex even if they have hits closer to the beampipe than the vertex coordinate. However stricter d_0 requirements are imposed on these tracks to suppress fakes [57]. These lower quality tracks attached later in the vertexing process will be referred to as *attached tracks*.

The number of tracks that form a DV is referred to as $n_{\text{Tracks}}^{\text{DV}}$, and the visible invariant mass from the four-momenta of the tracks is referred to as m_{DV} [57].

As of Run 3, LRT can be run during the HLT trigger on specific regions of interest, however it is too expensive to be run on the entire tracker. This means that there is now the capability in Run 3 for the development of triggering on DV objects that was not possible during Run 2 [54].

Searches for new physics using high mass DVs are discussed in the following two chapters.

Chapter 3

The Displaced Vertex + Jets Search (DV+Jets)

“No no you’re in a hateful corner of ATLAS phase space”

— Ynyr Harris, PhD Particle Physics, ATLAS, University of Oxford

This chapter details the "DV+Jets" search, where new long-lived particles (LLPs) beyond the SM are searched for using displaced vertex signatures in events that are triggered by multiple energetic jets. A combination of multiple-jet triggers were used, ranging from requiring four jets each with a requirement of $p_T > 100$ GeV to seven jets with a requirement of $p_T > 45$ GeV. The jets are only used as a trigger and there are no additional requirements that they originate from the LLP decay DV or not. This is to allow the analysis to be as general as possible, and inclusive to a wide variety of BSM models. Due to large low-mass and low-number of track DV backgrounds, the analysis can only be sensitive to models where the LLP decays produce DVs with a high mass and large number of tracks.

The term "prompt" or "promptly-decaying" will be used here to refer to anything that, to the detector, originates directly from the primary vertex (PV), i.e. originate from a particle produced in the proton-proton interaction that has a lifetime indistinguishable from zero.

Although the analysis is signature-driven, two benchmark models from R-parity violating (RPV) SUSY have been chosen to guide the analysis design as they tend to have final states with many jets. The first is an RPV scenario which involves the

SUSY strong sector, and will be hence referred to as the strong RPV gluino model, as shown in Figure 3.1a. Here pair-produced gluinos each promptly decay into quarks and a neutralino. These quarks become prompt jets. The neutralino is long-lived and then decays into a final state with only quarks. These quarks hadronize and are registered as displaced jets and vertices. Due to the prompt jets from the gluino decay, this process will have a large number of energetic jets in the final state, and so will have a high chance of passing a multi-jet trigger.

The second is also an RPV scenario, but involves new particles that are only from the SUSY electroweak (EWK) sector, henceforth referred to as the EWK RPV model, as shown in Figure 3.1b. Here, long-lived electroweakinos (which can be neutralinos or charginos) are pair-produced and each decays into a final state with only quarks. As with the strong RPV model, these quarks become displaced jets and vertices. However, the lack of prompt jets mean that this model will have a lower probability of passing a multi-jet trigger, which suggests the need for defining a separate signal region where the presence of a displaced jet requirement can allow for the loosening of the multi-jet requirements.

In both of these cases, the RPV coupling, λ'' , which allows the pair-produced SUSY particles to decay into a fully Standard Model final state, is very weak. This is due to this term violating baryon number which has not yet been observed in experiment, hence constraining this coupling to be very small [59, 60]. It is this weak coupling that makes that decay unlikely, therefore causing the electroweakinos to be long-lived.

This searches uses the full Run 2 proton-proton collision data set of 139 fb^{-1} at centre-of-mass energy $\sqrt{s} = 13 \text{ TeV}$. This is the first such Run 2 search in ATLAS. However, there was an equivalent search done in Run 1 data [61]. That search targeted slightly different models. The current search in this thesis improves upon the Run 1 search not only through a higher centre-of-mass energy and more data, but through an improved vertexing algorithm (VSI, as described in Section 2.3.4) and a more robust analysis design in terms of selections and background estimate method.

There are also a suite of other DV analyses being pursued in Run 2 ATLAS data, such as the early partial Run 2 DV+MET analysis (which uses missing transverse energy triggers) [2], a full Run 2 DV+muon analysis (which uses muon triggers) [57], and a full Run 2 DV+MET analysis which will be covered later in this thesis. Note that due to the nature of the analysis including jets, which will involve a much higher background from QCD multijet production, this needs to be done differently than

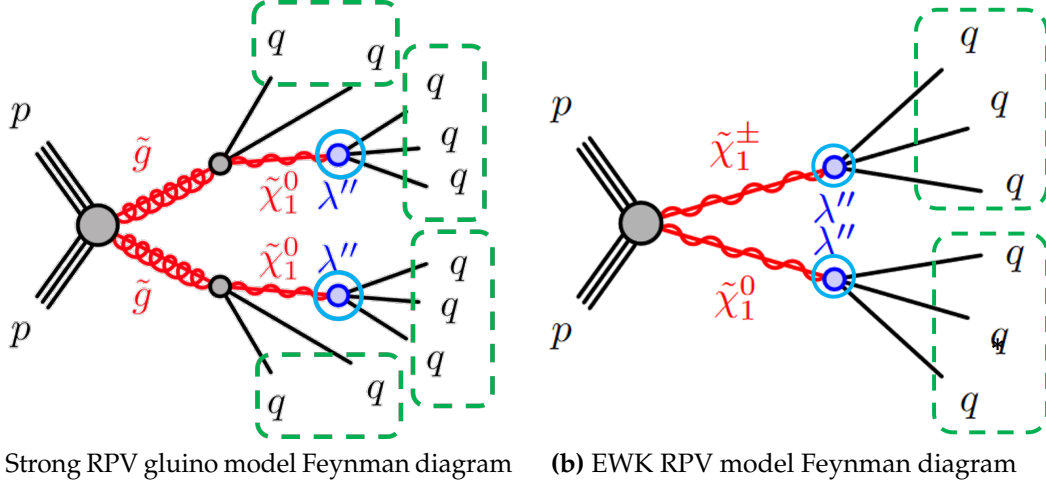


Figure 3.1: Benchmark models for the DV+Jets analysis. The blue circles represent the DVs, as well as where the RPV coupling λ'' is. The object surrounded by the green dots represent objects that the analysis triggers on, in this case multijets.

from the "cleaner" DV+muon and DV+MET analyses, and brings new challenges that did not have to be tackled by the Run 1 DV+Jets analysis.

My work on the analysis was primarily truth studies to understand our backgrounds, as detailed in Section 3.3, and the hadronic interactions background estimate, as detailed in Section 3.4.1. I was also responsible for the initial work on adapting the limit setting code from HistFitter from the DV+muon analyses to pyHF for this analysis, to produce plots as in Section 3.7, and studies on the combination portion of the combined background estimate. The limit setting code and combination portions were later passed over to other people in the analysis. I also contributed to other portions of the analysis, such as some studies on fixing the angles of the accidental crossings estimate, detailed in Section 3.4.2, and some studies in improving the selections which did not end up included in the analysis [62]. I also worked on the development of VBF Higgs Portal MC samples, which were later extended by Anna Mullin to full range of Higgs production modes, such that a Higgs Portal interpretation of this analysis could be performed. I also worked on finding and recommending fixes for some issues with the material map veto, detailed in Section 3.2.3. Although a temporary fix was implemented for this analysis, I developed a new veto system from the maps to address these issues that is later used in the DV+MET analysis, as detailed in Section 4.2

3.1 Data and MC Samples

The full Run 2 dataset of 139 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13 \text{ TeV}$ was used for this analysis. The data directly analysed in this analysis is collected by a set of multi-jet triggers as described in Section 3.2.1, however data collected from other triggers are also used in the analysis for background estimation as described in Section 3.5. The standard reconstruction for the particular software release used at the time of this analysis ("Release 21") did not include the large-radius tracking required to reconstruct displaced vertices (as described in Section 2.3.4). A filter was derived by the RPV/LL group to select events in data that would be interesting for several LLP searches to undergo bulk reconstruction. The filter included jet requirements that are more loose than the jet requirements defined in the analysis, in Section 3.2.1.

Simulated events for the signal models are needed in order to estimate yields, needed to both design signal regions and draw conclusions about the results. The strong RPV gluino model was simulated at leading order precision with the MADGRAPH5_aMC@NLO [2.6.2] generator for the hard-scatter matrix element calculation [63], then passed to Pythia 8.240 for showering and hadronization [64]. The pp interaction produces a pair of gluinos (\tilde{g}). Each gluino then decays into a long-lived neutralino-1 $\tilde{\chi}_1^0$ and quark-antiquark pair $q\bar{q}$. The long-lived neutralino-1 then decays into 3 quarks or antiquarks through the RPV coupling λ'' . All other RPV couplings are set to 0. The free parameters of the signal grid are therefore the gluino mass $m_{\tilde{g}} \in [1600, 2600] \text{ GeV}$, the LLP neutralino mass $m_{\tilde{\chi}_1^0} \in [10, 2550] \text{ GeV}$, and the LLP neutralino lifetime $\tau_{\tilde{\chi}_1^0} \in [0.01, 10] \text{ ns}$.

The EWK RPV model was also simulated at leading order precision with the MADGRAPH5_aMC@NLO [2.9.2] generator for the hard-scatter matrix element calculation [63], then passed to Pythia 8.245 for showering and hadronization [64]. The pp interaction produces two long-lived neutralinos or charginos that are in pure higgsino states. Possible combinations are $\tilde{\chi}_2^0 \tilde{\chi}_1^0$, $\tilde{\chi}_1^\pm \tilde{\chi}_1^\pm$, $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$, and $\tilde{\chi}_1^\pm \tilde{\chi}_1^0$. The neutralino and chargino masses are set to be degenerate, with the exception of the lightest neutralino being 1 GeV lighter than the others. This is to ensure that when these all decay, that they do all eventually decay into 3 quarks or antiquarks through the same RPV coupling λ'' . All other RPV couplings are set to 0. The free parameters of the signal grid are therefore the LLP neutralino mass $m_{\tilde{\chi}_1^0} \in [100, 1770] \text{ GeV}$, and the LLP neutralino lifetime $\tau_{\tilde{\chi}_1^0} \in [0.01, 10] \text{ ns}$.

Although the background estimate is largely data-driven, simulated events are still needed to study and understand the backgrounds, as well as to validate the background estimate methods. The dominant background for such a multi-jet trigger analysis will be dijet events. Therefore, dijet events were generated using Pythia 8.230 at leading order [64]. The samples are split up into slices corresponding to the p_T of the leading outgoing parton. Due to the computational cost of putting events through LRT reconstruction, only the slice corresponding to the $p_T \in [400, 800]$ GeV were used, as this slice is observed to largely overlap with data.

The effects of pileup, which are additional proton-proton collisions that occur within the same events that are not the hard-scatter process, are generated with Pythia 8.210 [64]. The pileup is then overlaid on top of the samples in three different profiles that correspond to different data-taking periods: MC16A (2015-2016), MC16D (2017), and MC16E (2018). Each signal sample grid point had 10,000 events were produced for each pileup profile. For the background sample, 1 million events were generated for MC16A and MC16D, while 13 million were generated for MC16E. All MC samples are propagated through a simulation of the ATLAS detector [65] using GEANT4 [66], after which the samples are processed in the same manner as data.

3.2 Selections and Signal Regions

There are several layers of selections that need to be performed in the analysis. Selections are applied to events, which include having jets that pass certain requirements and having at least one DV that passes specific selections, where selections are in turn also applied to the tracks of those DVs. If an event has more than one DV that passes selections, only the DV with the highest mass will be considered for the analysis. Therefore every event that is considered in the analysis will have exactly one DV. This is primarily to simplify the analysis as needing to take into account the possibility of multiple DVs per event and their correlation in background estimation, studies, and limit setting would take significant effort for relatively little gain, as virtually no events have more than one SR DV.

The selections were optimized where possible to reduce the number of background events in the signal region to about one. The motivation for this is that background estimation methods involving DV signatures tend to have large uncertainties due to modelling limitations. Constraining the analysis to "zero-background" or as close

as possible allows for a very clear distinction between a result with evidence of a signal scenario and without, even if there is a large uncertainty in the background estimate. This also allows for the sensitivity to scale approximately with the luminosity, requiring only a handful events to claim discovery.

In order to reduce the number of background events down to one, the multi-jet p_T requirements need be very high, as will be described in Table 3.4. Although this is suitable for the strong RPV gluino model which has many jets, these strict requirements would significantly reduce sensitivity to the EWK RPV model. For this reason, an additional signal region was designed, where the presence of one or two "trackless" jets allow for the loosening of the multi-jet p_T requirements. "Trackless" jets are jets that lack the presence of tracks reconstructed with the standard tracking algorithm ("standard tracks"), which is optimized for prompt track reconstruction. Their precise definition will be outlined in Section 3.2.1. These trackless jets likely correspond to a DV object, where the tracks of the DV get reconstructed in the calorimeter as a jet, however they are not reconstructed with standard tracking in the inner detector as they are not prompt. As the presence of a trackless jet implies that the event is more likely to contain a signal-like DV signature, this motivates allowing the strict jet requirements to be loosened, such that signal yields can be improved for the EWK model without excluding less background. This signal region will be referred to as the **"Trackless SR"**, as opposed to the nominal **"High p_T SR"** described earlier.

3.2.1 Event Selections

The event selections are detailed in Table 3.1. For an event to be included in the signal region, it needs to pass the multi-jet triggers in order for it to even be written to disk. A combination of triggers were used with various p_T requirements that depended on jet multiplicity. In general the p_T requirement decreases as the number of jets required increases. For example, there must be at least four jets in the event that each pass $p_T > 100$ GeV. However for a higher multiplicity trigger, there would instead must be at least seven jets in the event that pass $p_T > 45$ GeV.

Events then need to pass one of the RPV/LL filters in order to have been re-reconstructed with LRT. Offline jet multiplicity- p_T requirements were implemented to further reduce background as described in Section 3.2.1. Finally, the event must contain at least one DV that satisfies the DV selections described in Section 3.2.2.

Event Selections:

- Passes multi-jet triggers
- [Passes RPV/LL event filter requirements](the offline jet multiplicity- p_T requirements in the bullet point below are stricter than those of the filter)
- Passes either **High p_T** or **Trackless** SR offline jet multiplicity- p_T requirements [Section 3.2.1]
- Contains at least one DV that passes the **DV selections** [Section 3.2.2]

Table 3.1: DV+Jets Event Selections**Jet Selections**

The RPV/LL reprocessing filter is comprised of several filters, and includes an analogous region to both the High p_T SR and Trackless SR. The name RPV/LL originates from the SUSY RPV/LL group, which comprises several analyses, many of which target long-lived particles. The filter was designed by the subgroup before the analyses were started to include many different possible regions of interest for reconstruction. The RPV/LL filter requirements relevant to the High p_T and Trackless SRs are detailed in Table 3.2 and Table 3.3 respectively.

As these regions target signatures involving a large number of relatively high- p_T jets, a series of requirements are made where the event would pass if there is a certain number of jets that are all above a certain jet p_T threshold. The requirement on the jet p_T can be loosened as the jet multiplicity increases, as there is more of a guarantee that the configuration is signal-like if there are more jets reconstructed. These p_T requirements can be significantly loosened in the presence of a trackless jet, as the presence of such a jet with no standard-reconstructed tracks is likely caused by a DV, which increases confidence that the event is signal-like while significantly reducing background.

Note that all jets considered in this analysis are required to pass overlap removal and the basic *BadLoose* jet cleaning as described in [67], however without cuts on the charged-fraction of the jet or the fraction of jet energy that is deposited in the electromagnetic calorimeter.

Jets are defined as trackless if the sum of the scalar p_T of all tracks (which must have $p_T = 0.5$ GeV to be counted) associated with the jet is less than 5 GeV. The reason for this small amount of "prompt" tracks being permitted is due to pile-up.

High p_T RPV/LL Filter:

- 4+ jets with $p_T > 220$ GeV
- **OR** 5+ jets with $p_T > 170$ GeV
- **OR** 6+ jets with $p_T > 100$ GeV
- **OR** 7+ jets with $p_T > 75$ GeV

Table 3.2: RPV/LL Filter Jet Requirements for the High p_T SR.

Trackless RPV/LL Filter:

- 1+ **trackless** jets with $p_T > 70$ GeV and $|\eta| < 2.5$
- **OR** 2+ **trackless** jets with $p_T > 50$ GeV and $|\eta| < 2.5$

—AND—

- 4+ jets with $p_T > 120$ GeV
- **OR** 5+ jets with $p_T > 85$ GeV
- **OR** 6+ jets with $p_T > 70$ GeV
- **OR** 7+ jets with $p_T > 45$ GeV

Table 3.3: RPV/LL Filter Jet Requirements for the Trackless SR.

However, these multiplicity- p_T requirements are not strict enough to sufficiently lower the amount of background events in the SR. The requirements are therefore tightened in the offline multiplicity- p_T requirements, described for the High p_T and Trackless SRs in Table 3.4 and Table 3.5 respectively. All of these requirements are more strict than the RPV/LL reprocessing filter requirements, designed to only leave approximately one event of background in the SR each, while still having a suitable signal yield for each SR's targeted model. In order to keep the SRs orthogonal, events in the Trackless SR must fail the High p_T SR.

High p_T SR:

- 4+ jets with $p_T > 250$ GeV
- **OR** 5+ jets with $p_T > 195$ GeV
- **OR** 6+ jets with $p_T > 116$ GeV
- **OR** 7+ jets with $p_T > 90$ GeV

Table 3.4: Final Offline Jet Requirements for the High p_T SR.**Trackless SR:**

- **Fails High p_T SR selection**

—AND—

- 1+ **trackless** jets with $p_T > 78$ GeV and $|\eta| < 2.5$
- **OR** 2+ **trackless** jets with $p_T > 56$ GeV and $|\eta| < 2.5$

—AND—

- 4+ jets with $p_T > 137$ GeV
- **OR** 5+ jets with $p_T > 101$ GeV
- **OR** 6+ jets with $p_T > 83$ GeV
- **OR** 7+ jets with $p_T > 55$ GeV

Table 3.5: Final Offline Jet Requirements for the Trackless SR.

3.2.2 Displaced Vertex Selections

The DV selections are summarised in Table 3.6.

Only DVs whose vertex position falls within the fiducial region of $r_{DV} < 300$ mm and $|z_{DV}| < 300$ mm are considered. This corresponds to being contained before the end of the z of the pixel barrel, and before the r_{xy} of the first layer of the SCT [43].

Standard Model mesons containing heavy quarks, such as B -mesons, tend to live long enough to travel a few mm through the detector before decaying. When

these mesons are produced promptly in the pp -interaction, this is a Standard Model background for the analysis as they will form DVs with a few mm of displacement. In order to remove this background, a requirement on the radial position of the DV of $r_{xy,DV} - r_{xy,PV} > 4\text{ mm}$ is imposed. The vertex must also pass a goodness of fit requirement of $\chi^2/\text{DOF} < 5$. This χ^2/DOF quantity for the DV is determined internally during the running of the VSI reconstruction algorithm. This χ^2 test is done by checking if the tracks are compatible with the vertex location based on the d_0 and z_0 distance of each track to the vertex location, taking into account the covariance matrix representing the uncertainties of the track parameters.

As described in Section 2.3.4, standard VSI vertexing is done in two steps. Firstly, high quality tracks are merged through a seeding process. These high quality tracks must have $d_0 > 2\text{ mm}$, detector hits only at higher radii than that of the DV, and a momentum that is consistent with the PV-DV vector [58]. A second step is then run where lower quality tracks can be attached to a vertex. These lower quality tracks have loosened requirements, only requiring that the d_0 and z_0 significances be < 5.0 and a $p_T > 1\text{ GeV}$ [58]. The high quality tracks from the seeding process will be referred to as "core" tracks. Any lower quality tracks attached to the vertex in the second step will be referred to as "attached" tracks. To remove fake vertices, a requirement that the DV to have at least 2 core tracks that pass track cleanings is imposed. Note that both standard and LRT reconstructed tracks are given as inputs to the vertexing algorithm.

The above is what will be referred to as the vertex "pre-selection", which will be applied to all DVs that are studied including those in control regions, as these are more cleanings rather than selections. Additional selections are applied to define the signal region, which will be referred to as the "full selection".

The background caused by hadronic interactions is the highest overall. However, it dominates in low mass, low number of tracks, and if the DV lies inside detector material. Therefore, the DV selections were created such that a material map veto, as well as requirements that the DV invariant mass $m_{DV} > 10\text{ GeV}$ and that the DV contains five or more tracks (both of these quantities are calculated after track cleanings as described in Section 3.2.4) to remove as much as possible of this background. These requirements also ensure that no Standard Model decays end up in the signal region. The material map veto requirement uses a map of detector material to determine whether or not the position of the DV lies inside or outside of detector material. The signal region requires the DV lie outside of detector material. Details on the veto are described in Section 3.2.3.

A blinded validation region is chosen to be similar enough to the signal region but without containing significant signal in order to validate whether or not the background estimate is accurate. This is done outside of detector material, with the DV containing 4 tracks, and with DV invariant mass $m_{DV} \geq 20 \text{ GeV}$. This region will remain blinded until the final validation before unblinding in order to give additional confidence to the background estimate. This will therefore be referred to as the "blinded validation region" in order to distinguish it from other potential validation regions that may be accessed during the development of the background estimate. The blinded validation region was chosen to allow for the $m_{DV} \in [10 \text{ GeV}, 20 \text{ GeV}]$ range to be left as an unblinded validation region for the purposes of developing the background estimate methods. This is because this contains the transition region from where the hadronic interactions background dominates to where the accidental crossings background dominates

Any other region that passes the pre-selection that isn't the blinded validation region or signal regions are valid control or unblinded validation regions. Note that the blinded validation regions and signal regions in data are blinded until final validation and unblinding respectively. This means that during the development of the background estimate, the analysis team may not access information from events that fall in these regions. Only until there is significant confidence in the background estimates may the blinded validation region be accessed. If this is successful and the analysis obtains to approval to unblind the signal region, then the data in the signal region may be accessed. This is put in place in order to avoid human bias during the development of the background estimates. MC samples are not blinded. To better understand the post-pre-selection regions in the analysis in terms of m_{DV} , $n_{\text{Tracks}}^{\text{DV}}$, and inside/outside of material, they are shown laid out in a schematic in Figure 3.2.

3.2.3 Material Map Veto

There are a few aspects to take into account when designing the material map vetoes. To clarify terminology, the veto is the logic that makes use of a map to make a decision on whether or not a DV should be considered inside or outside of material.

Firstly, the detector geometry in real-life may be slightly different than that in simulation. For example, the pixel detectors in real-life are slightly offset from centre by less than a mm, and the description of non-active volumes of the detector in such as

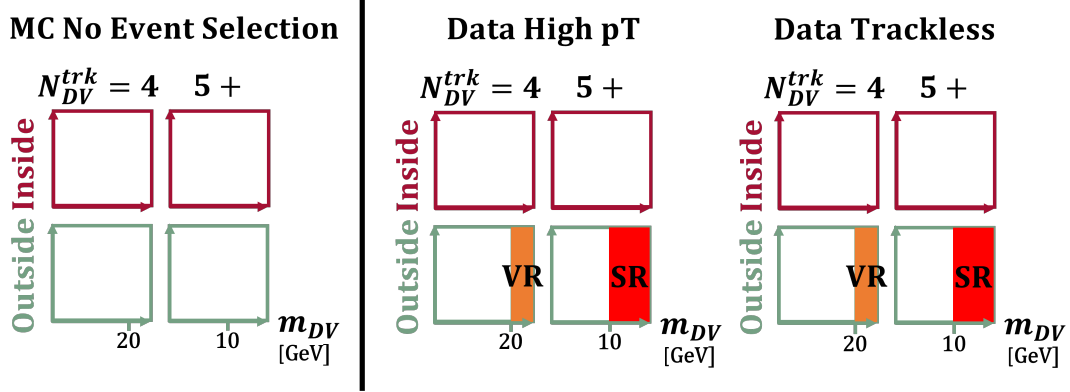


Figure 3.2: Schematic showing the different regions of the analysis in terms of m_{DV} , n_{Tracks}^{DV} , and inside/outside of material, for the MC with no event selection, as well as for the high p_T and trackless data regions. VR represents the validation region, SR represents the signal region.

Pre-selection:

- $r_{DV} < 300$ mm and $|z_{DV}| < 300$ mm
- $r_{xy,DV} - r_{xy,PV} > 4$ mm
- $\chi^2/\text{DOF} < 5$
- $N_{trks,core} \geq 2$, after track cleanings [Section 3.2.4]

Full selection:

- DV lies outside detector material as per material map veto (Section 3.2.3)
- $N_{trks} \geq 5$, after track cleanings [Section 3.2.4]
- $m_{DV} \geq 10$ GeV, after track cleanings [Section 3.2.4]

Table 3.6: DV+Jets DV Selections

cable bundles may be approximated in simulation. These changes are not significant enough for most analyses, hence why central MC production have not taken these effects into account. For this reason it is preferable to have a data-derived material map when dealing with data.

In order to do so, one must note that the geometry of the fiducial region comprises two primary sections: the inner section ($R_{xy} < 150$ mm) and the outer section ($R_{xy} > 150$ mm). The inner section comprises the actual pixel detector components and is fixed to the beampipe. The outer section comprises the SCT support structure which is ultimately fixed to the ATLAS cavern. Since the ATLAS cavern is used as the true coordinate reference, the beampipe may be slightly shifted, and the various pixel layers may also be slightly shifted off centre. For this reason the inner map is further split into 5 radial regions that each have an (x, y) offset defined.

In the inner section, the material map is determined by mapping out the density of vertices with low mass and low number of tracks that are not from K_s decays. The volume is binned and a threshold is applied such that if a bin contains a sufficient number of such vertices, the bin will be counted in the map as inside material. The 5 different inner section radial regions are additionally folded in ϕ to increase statistics. They are folded by a different factor based on the symmetry of the particular pixel layer in the region.

The outer region is too sparse to derive a map from this density method. Instead, a map taken from the detector geometry used for MC simulation is used as an initial guide, with certain regions that are not seen in data being removed and regions seen in data and not seen in the map being added [57].

Using the data-derived map directly as a veto, vetoing any DVs deemed by the map as inside material, is referred to as the "loose" data veto. Regions of the detector that would be vetoed by the loose data veto are shown in red for the $z = 0$ mm slice in Figure 3.3a.

Since the hadronic interactions background that dominates inside material does so by several orders of magnitude, it is optimal to be safe by not only checking the bin the DV coordinate falls in, but the immediate neighbouring bins in (R_{xy}, ϕ, z) as well. For this reason, the "strict" data veto was developed which considers a DV to be inside material if the bin the coordinate falls in or any of the immediate neighbouring bins is designated by the data map as containing material. Figure 3.3b shows in red DV

coordinates that would be vetoed as being inside material according to the data strict veto, at the $z = 0$ mm slice through the detector.

The reason why there is both a strict and loose veto is because in the loose veto there is background contamination in the outside material region and in the strict veto there is sufficient signal contamination in the inside region. Therefore, in data, studies looking "inside" material use the loose veto to obtain a pure background sample, where studies looking "outside" material use the strict veto to obtain a purer signal sample.

Of course when doing studies in MC, this data-derived map would be suboptimal, as locations of components are known to be shifted from the detector geometry that was simulated. For this reason, an MC map was created directly from the GEANT4 geometry. This MC map is then binned in (R_{xy}, ϕ, z) , such that the MC veto is defined by simply checking whether or not the coordinate bin the DV falls in is deemed by the MC map to contain material or not. Figure 3.3c shows, in red, DV coordinates that would be vetoed as being inside material according to the MC veto, at the $z = 0$ mm slice through the detector.

However, near the end of the analysis it was discovered simultaneously by myself and other members of the team that the MC map had designated pipes in the pixel layers containing cooling liquid as "outside" material, meaning the MC veto was incorrectly labelling a significant number of inside material DVs as outside material. This was an issue as it has severely affected background studies.

A permanent solution would be developed for the DV+MET analysis (Section 4.2), however for this analysis a temporary fix was used. As the data veto shift is quite small, the veto used for MC studies would be that the DV would be classified as inside/outside material only if both the old MC veto and the loose/strict data veto agreed the DV was inside/outside material. DVs with no such agreement were not considered for study.

Note that the data and MC maps were derived in a previous analysis [57] and is used across the SUSY DV group in ATLAS.

3.2.4 Track Cleanings

The track cleanings were primarily manually optimized by the team to reduce the number of background events to 1 for both signal regions. These are described




Figure 3.3: Visualizations of locations of DVs in $(R_{xy}, \phi, z = 0 \text{ mm})$ that would be considered by the different vetoes as inside material are denoted in red.

in Table 3.7. An upstream veto, the radius of the track's first hit in the detector, $R_{xy,hit}$, must be at the same or larger radius than that of the DV itself, $R_{xy,DV}$, is made to exclude fake tracks. No consideration is given to the uncertainty of the DV location in this veto. Various other requirements, such as on the track p_T , the impact parameter significance (" d_0 -sig") $|d_0|/\sigma(d_0)$, and the dot product angle $\Delta\alpha$ between the track and the $DV - PV$ vector, are made dependant on whether or not the track is core or attached, or the radial region of the DV. The angular requirements may, depending on the region, be a minimum angle requirement to remove prompt tracks, or a maximum angle requirement to remove pile-up tracks. The general radial regions used for track cleanings are $R_{xy,DV} \in [0, 25]$ mm ["BP"] (which represents inside of the beampipe), $R_{xy,DV} \in [25, 145]$ mm ["PIX"] (which represents the space), and $R_{xy,DV} \in [145, 300]$ mm ["OUT"] (which represents the space between the outer pixel layer and the SCT).

Track Cleanings

- All hits must $R_{xy,hit} \geq R_{xy,DV}$
- Track $p_T > 2$ GeV
- Track $p_T > 2$ GeV for **attached tracks** in "PIX" region
- Track $p_T > 4$ GeV for **attached tracks** in "OUT" region
- Track d_0 -sig > 10 in "BP" region
- Track d_0 -sig > 15 for **attached tracks** in "BP" & "PIX" regions
- Track d_0 -sig > 10 for **core tracks** in "OUT" region
- $\Delta\alpha(trk, PV - DV) > 0.2$ for any $p_T > 4$ GeV track in "PIX" & "OUT" regions
- $\Delta\alpha(trk, PV - DV) < \pi/2$ for **attached tracks**

Table 3.7: DV+Jets Track Cleanings

3.3 Sources of DV Backgrounds

MC simulations are not perfect and will inevitably mismodel certain aspects of the data. DV analyses are relatively non-standard and exotic in ATLAS, and so the ways MC may be mismodelling our backgrounds in our particular signal region is not a-priori understood. For this reason it is preferable that the background estimation method be as "data-driven" as possible, i.e. derived from the data rather than from MC. However, when designing such data-based methods, it is very important to understand what the methods are modelling, and by extension, what the sources of background are.

Although the MC cannot be relied upon to perfectly model the background, it can still be used as a guide for understanding in general what the backgrounds are and their general behaviours. MC can also be used for validation of our data-based methods through MC closure tests, where the background estimate methods are run on MC as if it was data and checked against the actual MC in the signal region. MC is also required to derive a correction and uncertainties for one of the background estimates, that of single-process DVs, due to very low statistics in both MC and data. This section details the relevant dedicated truth studies performed to understand, categorize, and catalogue the sources of DV backgrounds in the ATLAS ID. In this work, "truth" information refers to information saved about the "true" simulated particles that are being reconstructed in the MC that would not be accessible in data.

3.3.1 Framework for Truth Studies

In the MC production chain, first the main process of interest arising from a proton-proton interaction, known as the hard-scatter, is simulated at the level of elementary particles by a program known as a "generator". All possible Feynman diagrams for the process are generated at the requested order. For each collision event, the 4-vectors for all elementary particles involved are randomly drawn from their corresponding kinematic distributions, hence the name "Monte Carlo". The information for the 4-vectors of these elementary particles is stored in a "Les-Houches event" file format (LHE). Quarks hadronize after they are produced to form jets of bound hadrons, and so this is then simulated by a parton showering event generator, such as Pythia, which writes the information out into an EVNT file.

However the interactions of these particles with the actual detector has not yet been taken into account. GEANT4 takes each final state particle produced by the generator and then steps them through the detector geometry to model detector interactions such as hadronic interactions with material, energy loss, and the signal that is actually registered by the detector. GEANT4 also does this to any particle given by the generator that is long-lived enough to interact with detector material, the cut-off being defined at the generator level as particles with a $c\tau_0 > 1000$ mm ($\tau_0 > \approx 3.33 \times 10^{-9}$ s) [68]. The result of this is a HITS file which is identical in format to what is produced by the actual detector with real data, containing raw information on which elements of the detector registered a signal. The difference is that there is some level of truth-matching, where hits can be associated with the simulated truth-particle that created that hit.

The hits then need to be reconstructed into tracks, which would be our best estimates for the trajectory the real or truth particle took as it traversed the detector. It is at this stage that we need to implement the large radius tracking algorithm in order to reconstruct displaced tracks. Jets of quarks are also reconstructed as jet objects at this stage.

These files of basic objects, known as Analysis Object Data (AOD) files, are then processed to produce more group-specific higher-level objects in derivation files, known as DAOD files. We use DAOD_SUSY15, where the secondary vertexing algorithm is performed to group displaced tracks into displaced vertices that we can analyse. These files are not directly readable with regular programming languages such as bare Python or ROOT as they contain specialized links and objects that require the ATLAS software framework to interpret. As these files are relatively inaccessible and are not suitable for analysis, analyses run this through dedicated ATLAS software framework based programs that are typically shared across several analyses in a group (in this analysis it is FactoryTools). This specializes in reading these files and writing out data variables relevant to a particular analysis into a normal readable ROOT file known as an NTUP. The analyses then would use the more readily readable NTUPs for the actual data analysis in the analysis.

The information provided by the NTUP for the analysis, although typically sufficient for most purposes, was quite limited. Reconstruction in simulation is done the same way as in data, forming tracks from hits. However, these reconstructed tracks still need to be matched to the original simulated particle (if the track was presumably that particle) in a process called truth matching. The information about the truth particle in the simulation that is associated with a particular track is referred to as

truth information. Initially, the only information that was accessible in terms of truth information were the identifying barcode (unique number assigned to a particular truth particle in the simulation) and PDGID (a code corresponding to the identity of the particle, for example a π^+ meson would have PDGID +211 [10]) that was truth-matched to a particular track, the barcode and PDGID of a random direct parent of that track (as only one was recorded), and the truth-matching probability of the truth information to the track, which will be elaborated on later in this section.

In the earlier stages of the DV+Jets analysis, several inconsistencies were observed in MC between the individual background estimates and the truth distributions of what was being estimated, as originally defined using the truth information available in the NTUP. It was clear that the understanding of the backgrounds at the time was incomplete and risked not accounting for sources which could result in a false excess. In particular, the estimates of the independent backgrounds failed to add up to the expected number of background in closure tests. Other phenomena were additionally noticed such as the mass distributions of hadronic interactions unexpectedly differing inside and outside material and the shape of the accidental crossings estimate appearing different than expected.

Initiative was therefore taken to do a dedicated study to characterize, catalogue, and classify displaced vertex backgrounds in the ATLAS ID using the truth information available in the more inaccessible DAOD_SUSY15 files. The procedures and results of this dedicated study are outlined in the following.

The truth information available in the DAOD_SUSY15 files is still limited, particularly as there is some truth information that is simply not recorded anywhere by the simulation while running, specifically information involving pile-up particles. An effort was taken by the team to make more of this information available for the future DV+MET analysis. Note that there is also information that has been skimmed away from the previous stages of the MC production, in particular the fact that not all particles are actually passed down into the DAOD_SUSY15 file.

The relevant information provided in the DAOD_SUSY15 files after extraction are as follows. Not all truth particles in the event are included, however all those matched to tracks and all of their ancestors until the original proton-proton collision are included. Each of these particles has a barcode, PDGID, time and coordinates of production and of decay, 4-vector, as well as a list of the barcodes of their children particles and a list of the barcodes of their parent particles. This allows for the construction of a

decay tree and the tracing of the origin of a particular particle. Note that these are only provided for particles originating from the hard-scatter proton-proton interaction. No truth information is provided for tracks that originate from pile-up except for their truth-matching probability.

As for reconstructed tracks (reco-tracks), each track will, other than the 4-vector, have the associated truth barcode of the truth-particle it was matched to, along with the truth-matching probability. No "starting" or "ending" coordinate of the track is provided for reco-tracks, other than what can be inferred from the hit pattern. This can only be inferred from truth information.

If the track was from pileup or truth-matching was unsuccessful, the associated barcode is set to -9999 and no truth information can be provided. A barcode between 0 and 200,000 corresponds to a truth-particle that originated from the original generator. A barcode above and including 200,000 corresponds to a truth-particle that was created by GEANT4, either due to propagating it through the detector geometry or from some hadronic interaction with detector material.

The truth-matching probability (TMP) is a weighted measure of how many hits a reco-track has that are actually associated with the truth-particle in question. It is defined as

$$TMP = \frac{10N_{common}^{PIX} + 5N_{common}^{SCT} + N_{common}^{TRT}}{10N_{reco}^{PIX} + 5N_{reco}^{SCT} + N_{reco}^{TRT}}, \quad (3.1)$$

where $N_{common}^{PIX/SCT/TRT}$ represents the number of common hits in the PIX/SCT/TRT that the reconstructed track and the truth particle share, and $N_{reco}^{PIX/SCT/TRT}$ is the total number of reconstructed hits of the track in the PIX/SCT/TRT. The maximum of this ratio is 1, meaning all hits that form the reco-track are associated with the linked truth-particle. Truth-matching is done by calculating this probability for every particle which is associated with the hits of the track and choosing the one with the maximum truth-matching probability. In this work, "truth-matching probability" will refer solely to the maximum, which is associated with the truth-particle that is linked by the algorithm to the reco-track. If the truth-matching probability is low, it means that the reco-track is a fake track, reconstructed from hits that originated from many different particles [53]. Therefore, reco-tracks with $TMP < 0.5$ are deemed as fake tracks. Note that it should not be possible to have a TMP equal to or less than 0.

Reco-Track Type	TMP	Barcode
Gen	≥ 0.5	$> 0, < 200000$
G4	≥ 0.5	≥ 200000
True-PU	≥ 0.5	$= -9999$
Fake	< 0.5	Any

Table 3.8: This table summarises the definitions chosen to categorize different types of reco-tracks. Using these definitions, there are no reco-tracks that fall outside of these 4 categories.

It is important to note that many fake tracks may instead be due to a track travelling through the detector volume, undergoing some very collinear scattering with material and continuing on with essentially the same momenta as before, where only one of those two tracks is written as the truth track. If the the original track did not produce reliable hits, this case can result in a low TMP and the track being classified as a fake track, despite it essentially having the same properties as the track we want to study and it contributing to the DV properties as if it were a real track.

We can therefore choose to categorize the tracks using the above information into 4 main types, summarised in Table 3.8: those that originate from the generator (Gen), those that originate from GEANT4 (G4), pile-up tracks that are auxiliary tracks not associated with the proton-proton collision involved with the hard-scatter that is being studied in the event (True-PU), and spontaneous hits misreconstructed into a fake track (Fake).

The legend for the schematics shown in the rest of this chapter is shown in Figure 3.4.

MC Samples

Although the results of these studies are applicable to any analysis involving DVs in the ID, the context of this study was the DV+Jets analysis. Hence, the MC background samples used are those from the DV+Jets analysis, namely the dijet sample, as described in Section 3.1. Although the distributions of the background sources may change with different background samples, the actual categorizations of background are expected to be sample-independent. As mentioned in Section 3.1, pile-up conditions changed over the course of Run 2, and MC is produced in three different

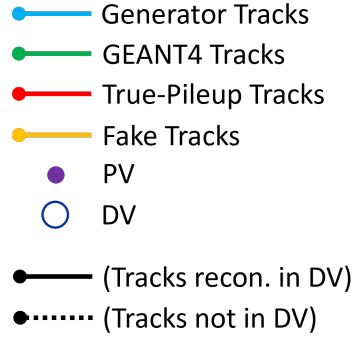


Figure 3.4: Legend of colours used to designate types of tracks in schematics in this following chapter. Solid lines represent tracks reconstructed in the DV, whereas dotted lines represent tracks not in the DV. A purple dot represents the PV, and a solid-line open circle would represents the DV.

campaigns to represent different conditions. Although event weighting does not matter for the purposes of cataloguing and categorization, it is important to note that for any plotted distribution in this section, the events are treated on an equal footing regardless of pileup campaign and not individually luminosity-compensated to their respective years or pileup conditions.

No event selections were used for these studies, as they would have simply lowered the already limited statistics and could hide sources of background. DV and track cleanings were applied as described in Section 3.2.2 and Section 3.2.4 respectively. Only the DV pre-selection was used, as defined in Table 3.6. Although this would not affect the overall categorization and cataloguing of DVs as done in this study, it is important to note that the process of applying such cleanings in the context of an analysis can cause a small number of DVs to change between categories. However, the general behaviours of the categories will still remain the same.

Defining the Originating SM LLP (OSMLLP) in Ancestry Studies

It was clear early on in the studies that the immediate truth-parents of reco-tracks do not provide information that is particularly meaningful for classifying DVs. Immediate parents may be short-lived particles created temporarily during a hadronic interaction that then decays into the particle associated with the track, as shown in the decay tree in Figure 3.5. Here we see a K^+ emerging from the PV and then undergoing a hadronic interaction with material creating many children. It has 4 immediate children that are reconstructed into the DV and appear to be detector-stable, except for one

which appears to undergo some sort of elastic scatter later in the detector. 1 additional track reconstructed in the DV does not directly originate from the K^+ , but rather from an intermediate $\eta'(958)$ particle. All of the tracks should be considered to be part of the same process. This means that one cannot use the immediate parent or even necessarily track truth origins as a good metric for whether or not a track belongs to the same process as the others.

There may be a secondary elastic hadronic interaction where the truth-particle associated with the track was actually not the particle created at the displaced vertex. GEANT4 may also occasionally change a particle barcode while in flight if dealing with scattering in the detector. Therefore, a novel method of meaningfully describing the origin of a reco-track in the context of DVs needed to be developed.

Ideally, any particle reconstructed in the detector can have their ancestry traced back to the originating Standard Model long-lived particle (OSMLLP) that was produced at the proton-proton collision and physically left and travelled away into space before either decaying displaced from the collision or being stable. The OSMLLP of a truth-particle can also be itself, for example as shown in the schematic in Figure 3.7. Such a definition would allow us to determine more meaningfully if tracks share a common displaced "origin" or not, as demonstrated in the schematic in Figure 3.7, as well as visible in the decay tree example in Figure 3.5. Note that since these OSMLLPs travelling out into the detector will no longer interact with each other due to distance, each particle in the detector geometry will have only a single OSMLLP.

However, some care must be placed in determining the exact threshold of how "displaced" is defined, specifically as there may be particles that live a non-zero amount of time but still decay insignificantly close to the proton-proton collision into long-lived particles that go on to create respective DVs.

Several candidates can be chosen for this radial threshold at which the particle must be produced before it and decay after it to be considered the OSMLLP of all of its descendants. For example $r_{xyz} = 0$ would be very strict and include any particles that have even an insignificant displacement. One could also choose $r_{xy} = 4$ mm, consistent with the analysis cut on displacement.

Figure 3.8 shows the distributions of particle production radii and decay times (in ct). In the lowermost leftmost corner in the underflow bin in both production time and decay time, are particles that are produced and decay immediately at the time of collision. These are mostly gluons and quarks from before the OSMLLP is made.




Figure 3.5: Decay tree of an event showing relevant truth particles. Edges represent particles, vertices are coloured by time distance from the PV (grey if they are at the PV itself). Particles highlighted in bold red are reconstructed in the DV. Dotted lines represent Gen tracks, solid lines represent G4 tracks. PU and fake tracks cannot be truth-matched to the hard scatter and are therefore not included. The particles are highlighted with the same colour share the same OSMLLP. Originating protons of the collision are out of the scope of the plot, but would be above. Note that not all particles involved in the event are included.




Figure 3.6: Truth particle event display of the event shown in Figure 3.5. All truth particles are shown, as well as the beampipe. The particles that have ended up reconstructed into the DV are outlined in red, where it is clear they all originate from the same interaction. Particles sharing the same OSMLLP as the reconstructed tracks are coloured in green. This appears to be a hadronic interaction with the beampipe.




Figure 3.7: Schematic of how the OSMLLP is determined. All tracks highlighted with the same colour share the same OSMLLP. The OSMLLPs are denoted in the schematic.

The mode of particles in the left lower quadrant represent short-lived particles that are produced and decay at essentially the PV. The mode of particles that appear in the left upper quadrant of the plot represent potential OSMLLPs: particles that are produced "promptly" and decay "displaced". Note that any particles in the upper overflow bin (above the red dashed line) are detector-stable. The mode of particles in the right upper quadrant are decay products of the OSMLLP: particles produced "displaced" and later decay even more "displaced". The particles on the diagonal line are extremely short-lived intermediate particles.

There appears to be a clear demarcation between the prompt and long-lived particle distributions at decay $ct \approx 1$ mm. The majority of the particles that live in the area just below $ct = 1$ mm but are not more clearly prompt are neutral pions. The cut on where the OSMLLP should be defined, in other words, the cyan lines, should be made to separate out these three modes: OSMLLPs, decay products of OSMLLPs, and extremely short-lived intermediate particles that essentially exist only at the PV. This was therefore chosen to be at $ct = 1$ mm, where the particle must be produced before this threshold and decay after this threshold to be considered the OSMLLP.

All of these candidate thresholds were tested after the classification study was performed, and it was found that no significant number of DVs changed category if the cut-off requirement was changed.

Therefore, for the purposes of these studies and the analysis:

The OSMLLP of a truth-particle is the particle's ancestor (not necessarily immediate, and can also be the particle in question itself), which was produced at $ct < 1$ mm and decayed at $ct \geq 1$ mm (or is stable). Note that promptly-decaying particles will not have a defined OSMLLP.

3.3.2 General Breakdown of DV Backgrounds

The general procedure for cataloguing and categorization involved sampling events and manually studying them to understand their properties and configuration. A programmable condition would be derived to describe this particular class of background and then events that pass this condition would be further sampled to verify if they indeed demonstrated similar behaviour or if there were sub-classes of DVs. This procedure was repeated iteratively until the only remaining classes of background were those with insignificant contributions.




Figure 3.8: 2D histogram of the production and decay times of all truth particles in a sample of events containing DVs with 3 or more tracks in a single DAOD file. The red dotted lines represent the underflow and overflow bins. The cyan lines represent the ct of 1 mm from the collision, where this appears to reasonably split the multimodal distribution into sensible regions.

In general, it was determined that one can roughly separate DVs into the following categories:

- **single-process DVs (tSP)**
- **single-process DVs with additional track (tAX for accidental crossings)**
- **other DVs (tOther).**

The "t" prefix denotes that these are categories in truth.

- A tSP DV is defined as a DV where all of its tracks share the same OSMLLP, and therefore, all originate from the same long-lived process.
- An tAX DV has all of its track except for one from the same OSMLLP.
- tOther DVs are for any other DVs that don't fall in any other category.

All of these categories each have their own subcategories which each exhibit their own unique behaviours which will be elaborated on in the following sections.

A plot of the breakdown of the DVs into these categories in terms of mass distribution, in this example inside and outside material for DVs with 4 tracks, is shown in Figure 3.9. tSP DVs dominate in the low mass region, whereas tAX dominate in the high mass region. tOther DVs are several orders of magnitude lower than any other source of background in any mass region, so they will be largely ignored.

3.3.3 Truth Single-Process DVs (tSP)

Truth single-process DVs are "true" DVs, in other words processes that are actually originating from the hard scatter and not the result of misreconstructing the DV.

An example of a decay tree of a tSP DV is in Figure 3.5, where all tracks in the DV originate from the same OSMLLP.

There are three major subcategories of tSPs with their own behaviours. **There are G4-only, Gen-Only, and G4+Gen tSP DVs**, defined by whether or not the tracks reconstructed in the DV consist of only those from GEANT4, only those from the generator, or both from GEANT4 and the generator respectively.




Figure 3.9: Breakdown of the mass distributions of the various truth-categorizations of DVs, here plotted for $n_{\text{Tracks}}^{\text{DV}}=4$ and inside material.

Gen-only tSP DVs represent SM decays, as they are true natural decays in flight at the generator level that do not involve an interaction with detector material, which would require GEANT4 to model.

G4-only tSP DVs largely represent hadronic interactions (HI). This means that some Standard Model long-lived or otherwise stable particle travelled through the detector and collided with detector material and underwent a hadronic interaction with the material to produce children tracks, which can only be modelled by GEANT4.

Note that HI can only be G4-only DVs and Gen-only DVs can only be SM decays. However it is possible some SM decays may be in the G4-only category if the particle had some scattering process in flight where it was taken over by GEANT4. However the G4-only category is dominated by HI.

G4+Gen tSP DVs represent an SM decay with a secondary hadronic interactions (HI). In order for all tracks to originate from the same OSMLLP and have both G4 and Gen children tracks, this means that there must have been an SM decay where one of the decay products shortly afterwards underwent a hadronic interaction with material.

A schematic of what these three categories look like is shown in Figure 3.10.




Figure 3.10: Schematic of the tSP sub-categories.

Figure 3.11 shows the shapes of the overall distributions inside and outside of material for the tSP category. The shapes of the overall distribution are different inside and outside material. The reason for this is because there are still a significant amount of SM decays, mainly from B mesons, that pass the displacement cleaning cut of 4 mm. They therefore contribute a significant fraction and may even dominate outside of material changing the shape. However, at higher masses hadronic interactions dominate, therefore given the signal region cut $m_{DV} \geq 10$ GeV, no SM decays make it into the signal region due to the sharply falling tails and no known SM hadrons of such high mass.

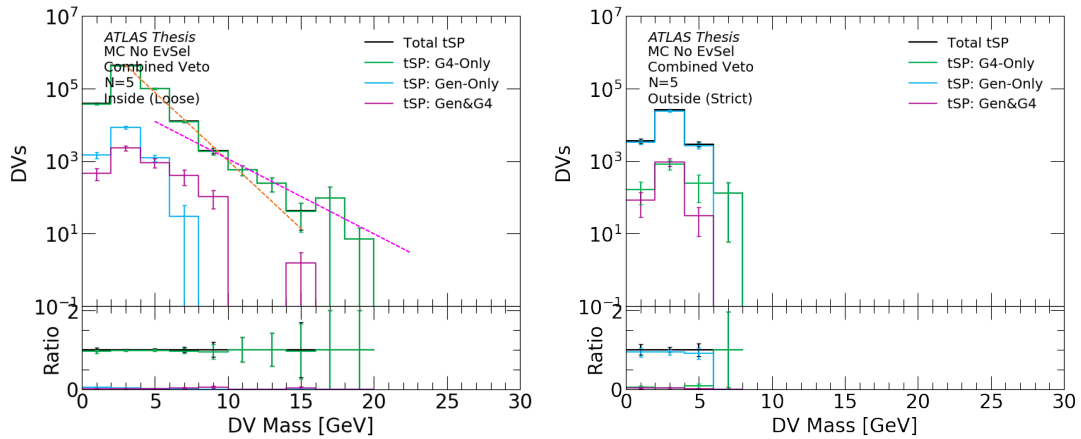


Figure 3.11: Breakdown of the mass distributions of the tSP DVs, plotted for $n_{Tracks}^{DV} = 5$ inside and outside of detector material. For the inside plot, the two dotted lines are shown in order to guide the eye to demonstrate the inflection in slope in the hadronic interactions curve, here occurring at about $m_{DV} = 10$ GeV.

Hadronic interactions have an additional complication. As visible in Figure 3.11 the "slope" of the exponential portion of the HI curve appears to have two components.

Ignoring the rise and peak, there is a steeper-sloped exponential at lower masses which then changes into a more shallower-sloped exponential at higher masses.

Upon closer inspection of these events, it appears that the population of DVs in the shallower high mass exponential tail have a different configuration to those in the lower mass sharper exponential tail despite both being HI.

As shown in the comparison and schematic in Figure 3.12, unlike in regular HI where the momentum is spread somewhat evenly amongst the children particles, these DVs in the shallower high mass exponential appear to have their mass extremely dominated by a single high- p_T track that is relatively collinear with the PV-DV vector. This is likely due to a unusually high momentum OSMLLPs originating from the PV, which would therefore allow a high momentum collinear outgoing track after the interaction.

There is however not a simple variable that conveniently and cleanly separates all of these "elastic collinear hadronic interactions" from regular hadronic interactions, and the need to develop one is unnecessary. Rather it seems that these two components are ends of a spectrum that co-exist across the mass distribution, and just one is more prominent in low mass and the other in high mass. However one should note the existence of this, as care would be required in the estimation methods to properly account for this behaviour.

3.3.4 Truth Single-Process Plus Track DVs (tAX)

These DVs consist of a "true" DV, but with an extra track crossing from a different process that has accidentally been added to the DV during its reconstruction.

We can divide up this category based on the identity of the crossing track: **Gen**, **G4**, **PU**, or **Fake**, with schematics shown in Figure 3.13.

These each have different shapes and behaviours in terms of their mass distributions as shown in Figure 3.14.

The Gen-tAX DVs have only a single peak at low mass, save for some noise in high mass. This is because Gen tracks that cross the DVs must originate close to the PV, which means in order to cross the DV such that it is reconstructed within it, it is restricted to be very close to the PV-DV vector. This means that it cannot add too much mass to the DV due to the relatively small momentum that is transverse to the

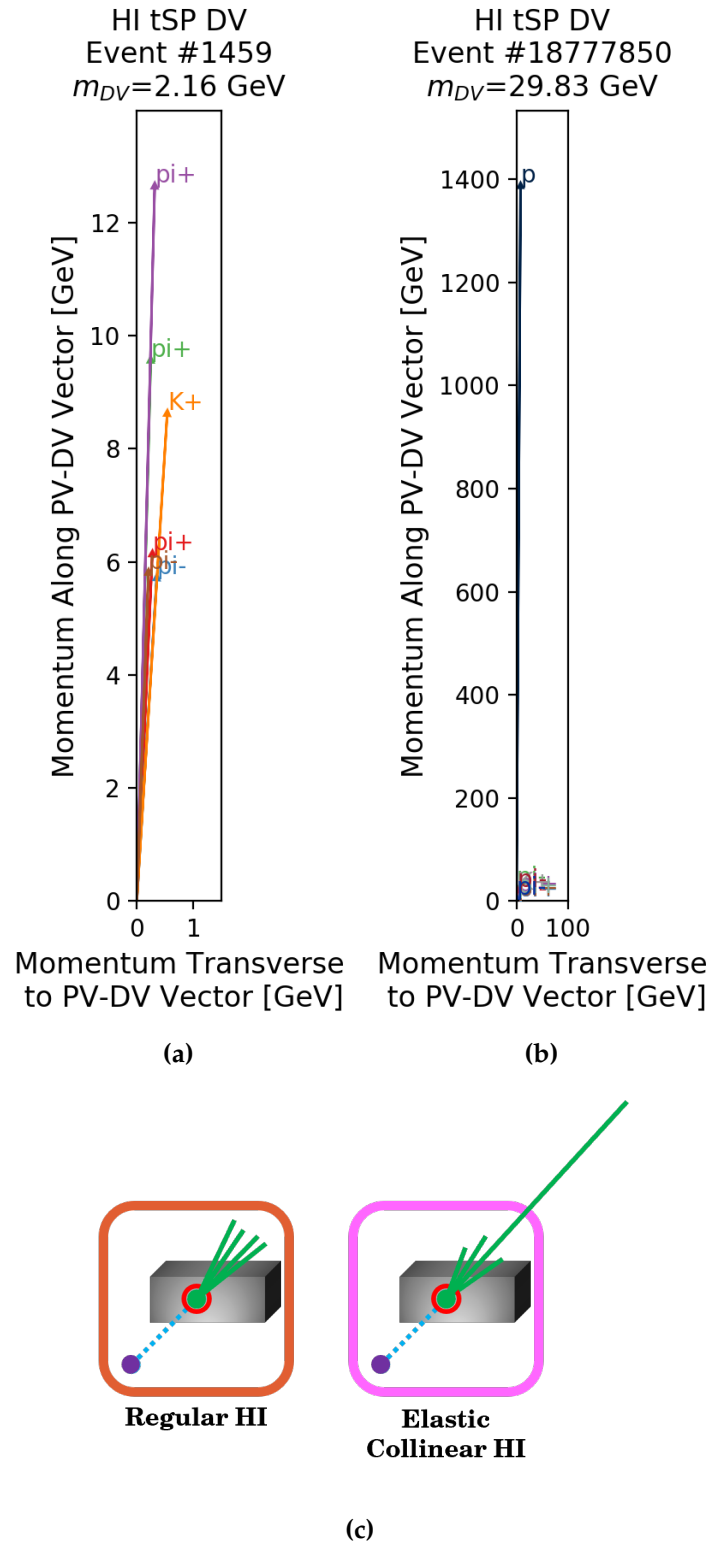


Figure 3.12: Example of an a) "regular HI" DV and a b) "elastic collinear HI" DV. These plots have the PV-DV vector aligned to the y-axis, whereas the absolute value of the momentum in the direction transverse to the PV-DV vector is plotted in the x-axis. These figures keep an equal aspect ratio in order to accurately reflect the scale of the relative orientations of the tracks. c) shows a schematic of these DVs.





Figure 3.13: Schematic of the tAX sub-categories.

Figure 3.14: Breakdown of the mass distributions of the tAX DVs, plotted for $n_{\text{Tracks}}^{\text{DV}}=4$ inside detector material.

PV-DV vector, and that the shape will largely resemble that of HI. The few events in high mass are due to small mass DVs with an unusually high p_T Gen track crossing.

The PU-tAX DVs have a slowly decreasing distribution over all mass. This is because the crossed track explicitly does not come from the same PV, and therefore could come from anywhere. This means that the restriction that it must be reconstructed in the DV does not limit the angle at which it crosses, which can be as small or as large as possible, only restricted by the length of the interaction region. This large continuous range is what causes the large continuous and relatively flat distribution, where the lower mass corresponds to more collimated crossing angles, and the high mass to larger crossing angles.

The G4-tAX DVs have a combination of behaviours of the Gen-tAX and PU-tAX, with a low mass peak followed by a shallow decreasing tail. This is because, being a G4 track, it is usually a child of a track originating from the PV that decays displaced, and therefore does not need to originate from the PV itself. This means it has an extra degree of freedom. The low mass peak corresponds to cases where the G4 track originates close to the PV, and the high mass tail corresponds to where the G4 track originated further from the PV and closer to the DV.

The Fake-tAX DVs have similar distributions to the Gen-tAX DVs. This hints that many of the tracks classified as fake originate from the PV. This is consistent with the observation that many reconstructed tracks classified as fake in fact have momenta consistent with their associated truth track despite that particular truth track not contributing enough hits to have a high enough TMP. This is in part because hits that contribute more to the TMP originate from the PIX, whereas hits that contribute more to the p_T originate from the TRT, meaning that a missed hit in the PIX can cause the TMP to lower significantly without greatly affecting the p_T .

It is clear that at high mass, the PU-tAX dominates with a large contribution from G4-tAX.

3.3.5 Other Categories of Truth DVs (tOther)

The rest of the categories can be separated into the following. A sub-DV here refers to a subset of the DV comprising tracks that all share the same OSMLLP. However in

terms of underlying mechanisms, we can generally classify these as **merged vertices, doubly accidentally crossed DVs, and other**. These will be described below.

DVs comprised of two sub-DVs merged together, both from the same hard-scatter (tHS-MV).

DVs comprised of a sub-DV from the hard scatter and PU tracks (tDVw2+PU). As there is no truth information with PU, it is impossible to know if all the PU tracks are from the same OSMLLP in their respective PV (in which case this would be a merged vertex process), or if the tracks were all from separate origins (in which case this would be a multiple accidental crossings), or some mix of the two.

DVs comprised of a sub-DV plus two tracks, where the tracks have different origins or OSMLLPs (t2xCertainAX). These are doubly crossed DVs. Both crossed tracks cannot be from PU at the same time, as we would not be able to tell the mechanism.

DVs comprising only PU-tracks (tPU-Only) It is impossible to know what the underlying mechanism of these DVs are and whether or not their tracks all come from the same PV (separate from the PV of the hard scatter), and even if so whether or not the underlying mechanism would count as a tSP, tAX, merged vertex, doubly-crossed DV, or something else.

Finally there are just miscellaneous DVs (tMisc) that do not fall into any of these categories, usually due to the tracks being from multiple different origins.

Schematics of the categories are shown in Figure 3.15. The mass distributions of these categories are shown in Figure 3.16. The most prominent contribution are from tHS-MV, with the second largest being from t2xCertainAX. There are limited statistics, however it is safe to assume the other sources of background are insignificant and do not need to be explicitly accounted for.

It is important to note, however, that given the information provided by the MC at truth-level, there remains a level of ambiguity and it is impossible to know precisely the exact mechanism of the DVs in some of these subcategories. This has been kept in mind for future analyses, where efforts are now underway during the DV+MET analysis to produce specific MC which keep truth information from PU tracks, in order to more accurately classify DVs.




Figure 3.15: Schematic of the tOther truth categories.




Figure 3.16: Breakdown of the mass distributions of the tOther DVs, plotted for $n_{\text{Tracks}}^{\text{DV}} = 5+$ inside detector material.

3.4 The Combined Background Estimate

The three major mechanisms behind the sources of background are "true" DVs (eSP), accidental crossings (eAX), and merged vertices (eMV). The combined background estimate aims to estimate these individually and combine them together to form an estimate. Other sources of background, such as DVs with two additional unrelated tracks or pileup-only DVs, are insignificant compared to the sizes of these primary background sources in all regions and so are therefore not considered.

How these estimated categories relate back to the truth categories of background as described in the previous section is non-trivial and would depend on the estimation method. In some cases it may not be possible to fully separate or map the estimated background directly to the truth backgrounds. This will be the case if the estimate, using only information available in the data, can only estimate a proxy of a category defined from the truth information. For example, the accidental crossings estimate adds a track onto existing DVs. However, some of the DVs the track is getting added to may already have a crossed track, as the method cannot tell if this is the case using information available only in data. Hence this estimate does not purely estimate the tAX category. Estimate categories are therefore prefixed by an "e" to avoid confusion with the truth categories prefixed by a "t".

As the MC cannot be relied upon for the background estimates as its mismodelling in this context is not well understood, the goal is to make the background estimates as data-driven as possible and to rely as little on the MC as possible. This however may create challenges for certain backgrounds due to the inherently few raw events available to study in the high mass regions in both data and MC, where it is essential we model properly due to the signal region also being in high mass.

This section will cover the individual estimation procedures, the closure test in MC, and the combination.

3.4.1 Estimating Single-Process DVs (eSP)

This eSP background can be mapped to the tSP category fairly easily. The goal is to properly estimate the number of single-process DVs, in other words "true" DVs that have not been affected by reconstruction effects.

In theory, it is possible to create an angular variable that could separate out SP from other backgrounds such as the maximum angular separation of the tracks. Unfortunately, since signal would tend to be SP-like, this causes issues with signal contamination as well as too few events to create a reliable estimate. Instead, after many attempts at an estimate, a functional fit method was developed.

Below $m_{\text{DV}} = 10 \text{ GeV}$, in DVs with any number of tracks, both inside and outside material, the tSP dominates all other backgrounds by at least an order of magnitude. It is only above this point where the other backgrounds become prominent. Therefore, one could assume that everything $m_{\text{DV}} \leq 10 \text{ GeV}$ is SP, and fit a function that follows the form to this low mass region to allow for extrapolation to the high mass region.

The initial fitting function is as follows, to be fitted in the $m_{\text{DV}} \leq 10 \text{ GeV}$ portion of the given region that is being estimated:

$$\rho(m_{\text{DV}}) = \frac{1}{\frac{1}{A(m-B)} + \frac{1}{e^{-\frac{m-C}{D}}}}. \quad (3.2)$$

Note that, this being the mass distribution, the actual number of DVs visible in a bin will be the integral of this function between the bin edges. This function intends to model the linear-like rise ($A(m-B)$) that peaks before undergoing an exponential-like decay ($e^{-\frac{m-C}{D}}$) that is the general shape of tSP. The reason for the unusual combination of the two functions via this reciprocal of a reciprocal sum form is to ensure a smooth transition between the density functions.

However, as mentioned in Section 3.3.3, there is a secondary exponential tail from elastic collinear HI that needs to be taken into account. Unfortunately, the inflection point in most regions occurs at too high mass to be able to properly fit a secondary exponential tail at low mass for extrapolation.

Since our signal region is high mass outside of material, we can use the high mass inside material regions as control regions for the secondary exponential and extrapolate to the signal region. However, in data the high mass region has contamination from other components of background. Although it is possible to roughly separate these out with an angular variable $((\Delta R)_{\text{max}}$, described later in the section), deriving the templates directly in data inside material and extrapolating to outside material would deprive us of a high statistics validation region for this corrected estimate. Although

we would be able to use the high mass outside material $n_{\text{Tracks}}^{\text{DV}}=4$ region between $m_{\text{DV}}=10$ GeV and 20 GeV for validation, this unfortunately does not have sufficient statistics as it is outside material, and it would be preferable for these SP DVs that dominate inside material to have a validation region inside of material.

Therefore, the correction to the tail will be derived from inside material high mass regions in MC, and then validated in an MC closure test as well as in inside material data with a $(\Delta R)_{\text{max}}$ cut.

To understand how we will derive this correction, for simplicity, assume we only consider the SP DVs mass distribution above about $m_{\text{DV}}=5$ GeV, where it is essentially only the exponential component. In other words, we can write

$$\rho_{\text{IN}, n_{\text{Tracks}}^{\text{DV}}, \text{MC No EvSel}}(m) \approx e^{-\frac{m-C}{D}} + e^{-\frac{m-C_2}{D_2}}, \quad (3.3)$$

where $e^{-\frac{m-C}{D}}$ represents the original exponential fitted at low mass, and $e^{-\frac{m-C_2}{D_2}}$ represents the secondary exponential that dominates at high mass. This secondary exponential can be obtained by fitting the full mass range of tSP DVs inside of material in the MC. Specifically, the original function is fitted in the inside material, $m_{\text{DV}} \leq 10$ GeV region of the given $n_{\text{Tracks}}^{\text{DV}}$ tSP MC with no event selections. Then the "corrected function",

$$\rho(m_{\text{DV}}) = \frac{1}{\frac{1}{A(m-B)} + \frac{1}{e^{-\frac{m-C}{D}}} + \frac{1}{e^{-\frac{m-C_2}{D_2}}}}, \quad (3.4)$$

is fitted, however the parameters already determined in the regular fit are fixed, and only C_2 and D_2 are actually fitted here.

If we try to go from the MC no event selection inside material region for a given $n_{\text{Tracks}}^{\text{DV}}$ to the equivalent region in data with event selections, the exponentials will be affected by not only the selections but by differences in the MC and data. However, as they both remain exponentials, only with potentially a change in slope and amplitude, this change must be parameterizable by the multiplication of a factor of the form $Se^{-\frac{m}{s}}$. If the selections and the difference between MC and data affects both components in the same way, this factor will be the same in both components.

One can attempt to verify this in a rough approximation by plotting the tSP component against the data with a cut on $(\Delta R)_{\text{max}}$ in order to select for tSP DVs. $(\Delta R)_{\text{max}}$ is

defined as the maximum opening angle between a DV track and the summed track momenta of all other tracks in the vertex. Accidentally crossed tracks will usually cross at a large angle to the rest of the tracks in the DV, allowing us to use this variable to help separate tSPs from the accidental crossings background, as shown in Figure 3.17. $(\Delta R)_{\max} < 0.8$ is therefore selected as the cut, and the data with this cut is shown against the tSP DV distribution in Figure 3.18, where relatively good agreement is seen.




Figure 3.17: $(\Delta R)_{\max}$ distribution of tSP and non tSP DVs in MC. A cut is selected at <0.8 to select for a high purity sample of tSP DVs.

As these match within uncertainties, one can write in the high mass region the approximation that

$$\rho_{\text{OUT}, n_{\text{Tracks}}^{\text{DV}}, \text{Data w. EvSel}}(m) \approx S e^{-\frac{m}{s}} \rho_{\text{IN}, n_{\text{Tracks}}^{\text{DV}}, \text{Data w. EvSel}}(m) \quad (3.5)$$

$$\approx S e^{-\frac{m}{s}} e^{-\frac{m-C}{D}} + S e^{-\frac{m}{s}} e^{-\frac{m-C_2}{D_2}} \quad (3.6)$$

where $\rho_{\text{OUT}, n_{\text{Tracks}}^{\text{DV}}, \text{Data w. EvSel}}(m)$ is the DV invariant mass (m) distribution with $n_{\text{Tracks}}^{\text{DV}}$ number of tracks in data with the event selections but outside of detector material, and $\rho_{\text{IN}, n_{\text{Tracks}}^{\text{DV}}, \text{Data w. EvSel}}(m)$ is the same but inside of detector material. C, D, C_2, D_2 are defined as in Equation 3.4, and S and s define the exponential factor that characterize the difference in the distribution when going from inside to outside of material.

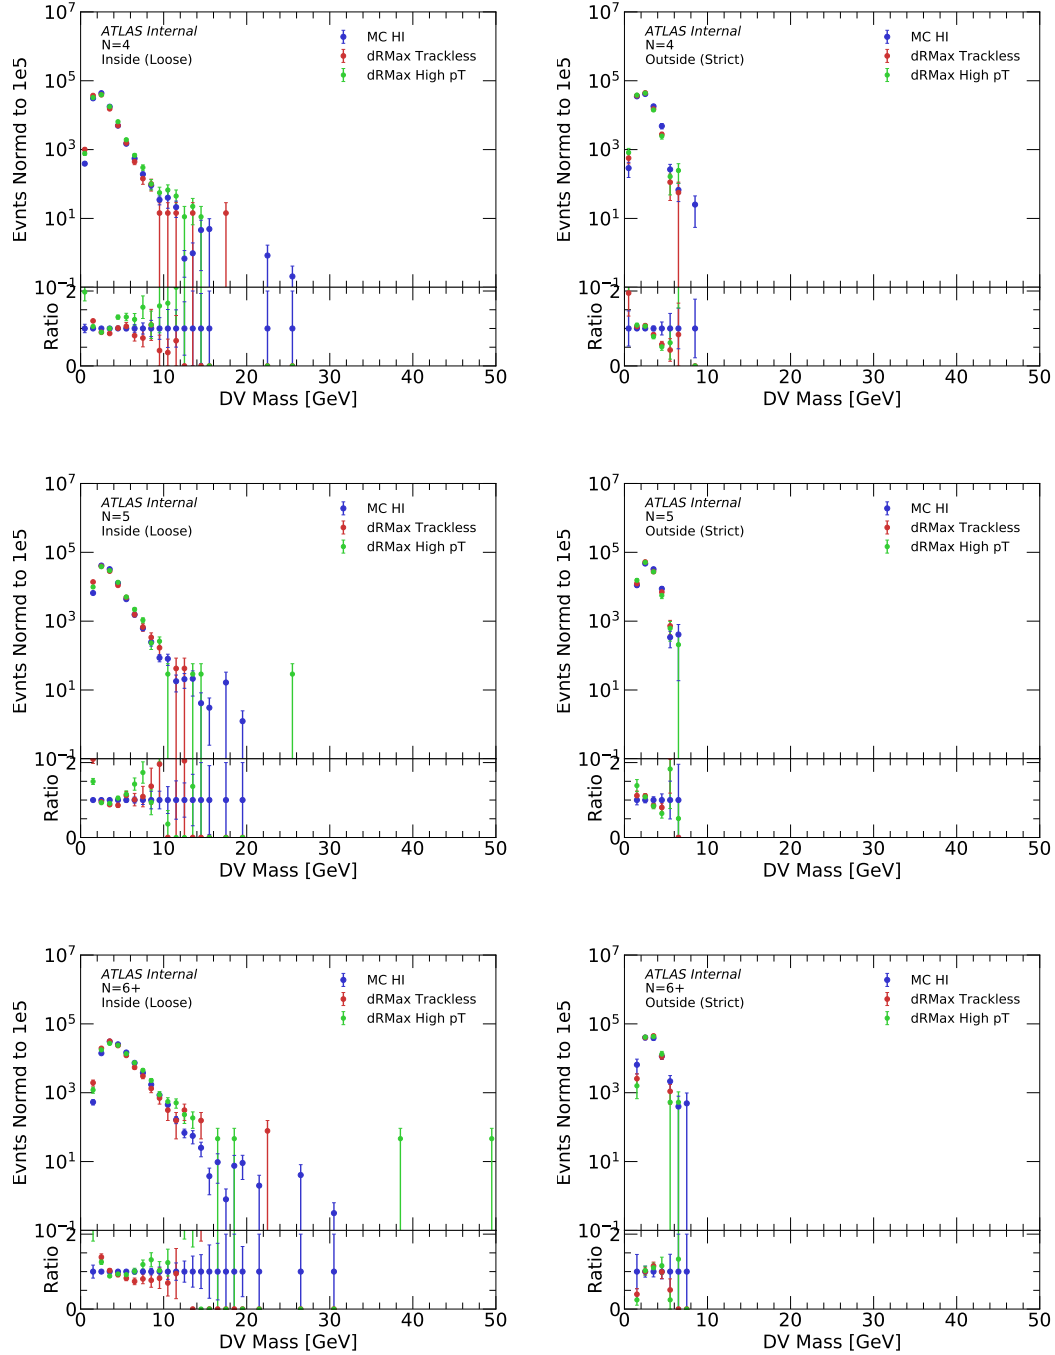


Figure 3.18: The mass distribution of $t\bar{t}SP$ MC is plotted alongside data with a $(\Delta R)_{\max}$ cut, in order to get a proxy for what the $t\bar{t}SP$ components of data with selections look like. The first column is for inside material and the second column is for outside material, with four (first), five (middle) or six or more tracks (last).

Going from inside to outside material, the high mass region should only be from true HI, and there should not be any significant difference in interaction between for example a metallic nucleus inside detector material and an air molecule outside detector material. Therefore, the only difference should be a change in amplitude and one can represent going from inside to outside by applying a factor A to the inside mass distribution to obtain the outside mass distribution. This is unfortunately very difficult to verify in MC as there are so few statistics in the outside region that there are typically too few tSP DVs in the region where the original exponential begins to dominate and the effect of the B-decay peak begins to vanish, as shown in Figure 3.19. However, it is important to note: the estimate is fundamentally limited by our knowledge about the behaviour of this secondary component, due to the lack of statistics available in the tSP MC. Therefore, as will be discussed in Section 3.4.1, not only will the regular fit uncertainties be propagated to the final template, but the statistical uncertainties in this tSP MC control region will be as well. However, one must appreciate that there is good agreement in the region where there are sufficient events to draw a conclusion.

These two assumptions allow the estimation inside or outside material of some given region from a control region of corresponding $n_{\text{Tracks}}^{\text{DV}}$ in MC with no event selections but containing only tSP DVs.

A schematic of the information available and a visual representation of the estimation method in terms of the equations is shown in Figure 3.20. To clarify, the primary function and the secondary exponential are both fitted in the tSP MC inside control region. The primary function is then fitted in the low mass region in data. The secondary exponential in data is then set to be the secondary exponential in the tSP MC control region but normalised by the ratio of the primary exponential in data over tSP MC.

In other words, if $\rho^{(1)}(m)$ was the primary exponential and $\rho^{(2)}(m)$ was the secondary,

$$\rho_{\text{OUT}, n_{\text{Tracks}}^{\text{DV}}, \text{Data w. EvSel}}^{(2)}(m) = A S e^{-\frac{m}{s}} e^{-\frac{m-C_2}{D_2}} \quad (3.7)$$

$$= \frac{\rho_{\text{OUT}, n_{\text{Tracks}}^{\text{DV}}, \text{Data w. EvSel}}^{(1)}(m)}{\rho_{\text{IN}, n_{\text{Tracks}}^{\text{DV}}, \text{MC No EvSel}}^{(1)}(m)} \rho_{\text{IN}, n_{\text{Tracks}}^{\text{DV}}, \text{MC No EvSel}}^{(2)}(m). \quad (3.8)$$

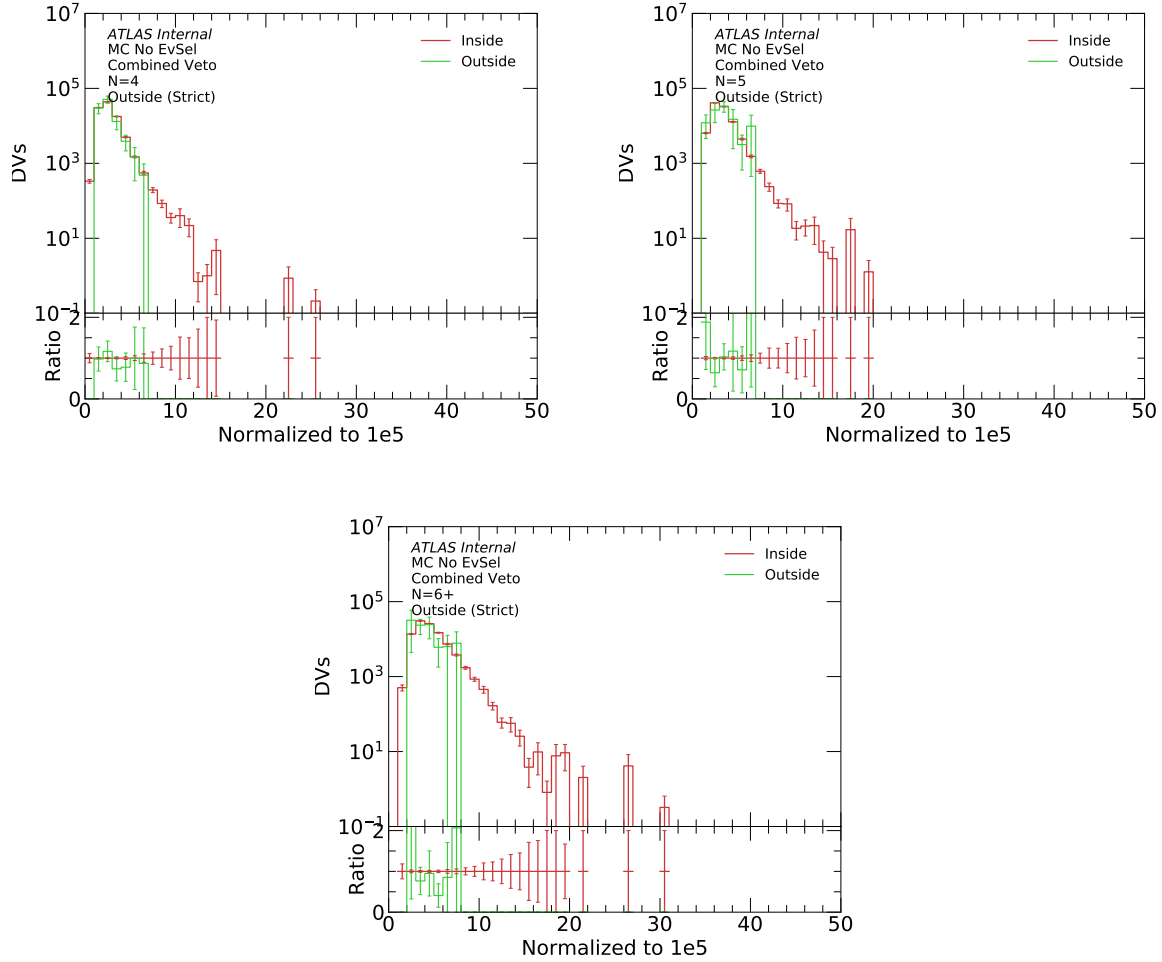


Figure 3.19: tSP MC DVs that only have GEANT4 tracks (in order to isolate the "true HI" DVs that dominate in the high mass from possible decays that only exist in low mass) compared inside and outside of material in order to determine whether or not their shapes are comparable. The DVs have four (left), five (middle) or six or more tracks (right).

	tSP MC, No Ev. Sels.	Data With Ev. Sels.
Inside	$e^{-\frac{m-C}{D}} + e^{-\frac{m-C_2}{D_2}}$	$[S_e \frac{x}{s_e} \frac{m-C}{D}] + S_e \frac{x}{s_e} \frac{m-C_2}{D_2}$
Outside	$[A_e \frac{m-C}{D}] + A_e \frac{m-C_2}{D_2}$	$[ASe \frac{x}{s_e} \frac{m-C}{D}] + ASe \frac{x}{s_e} \frac{m-C_2}{D_2}$

Fitted in Primary Fit | Fitted in Secondary Fit | Cannot be fitted | Could be fitted with dRMax cut

Figure 3.20: Schematic showing the equations of the mass distribution in terms of the primary and secondary exponentials in the different regions.

Uncertainties

With the low mass primary function fitting, a covariance matrix is provided by the fitting algorithm. In order to obtain the fitting uncertainty, sample distributions of the primary function are generated from sampling multi-variate gaussian distributions for each of the parameters using the corresponding elements of the covariance matrix. Sample distributions falling outside the range limits of the parameters, set to ensure proper physical behaviour of the function, are ignored. The standard deviation of these toys in a given mass bin is therefore written out to be the fitting uncertainty in that mass bin. Note that these uncertainties are correlated across bins, and so the uncertainty of any given mass range must be computed separately, and cannot be calculated from the uncertainties of individual bins.

As previously mentioned, the accuracy of the secondary function fit is primarily limited by our knowledge of its behaviour, which is limited due to the lack of events available in the tSP MC to study. Therefore, the statistical uncertainties in this tSP MC control region will be propagated to the template, normalised by the same ratio of the primary fits in the data and tSP MC control region.

Closure Test

Although the estimation method uses tSP MC, one can still do a limited closure test in the full MC to check if the primary fitting in the full MC works sufficiently well to model the tSP, as well as whether or not the ratio method for the secondary fit is functional. The full MC was therefore passed to the algorithm as if it was data. The results of the closure test are shown in Figure 3.21 for all MC regions. The estimate matches the true tSP MC to within uncertainties, so no additional non-closure uncertainty needs to be applied.

3.4.2 Estimating Accidental Crossings (eAX)

The general method for estimating accidental crossings predates the author's joining of the analysis team. The method involves estimating the number of $n_{\text{Tracks}}^{\text{DV}}$ -trk DVs with crossed tracks by taking $(n_{\text{Tracks}}^{\text{DV}}-1)$ -trk DVs and adding on an extra track, and then normalising the resulting mass template by a crossing factor representing how often crossings occur. The procedure is outlined in Figure 3.22.

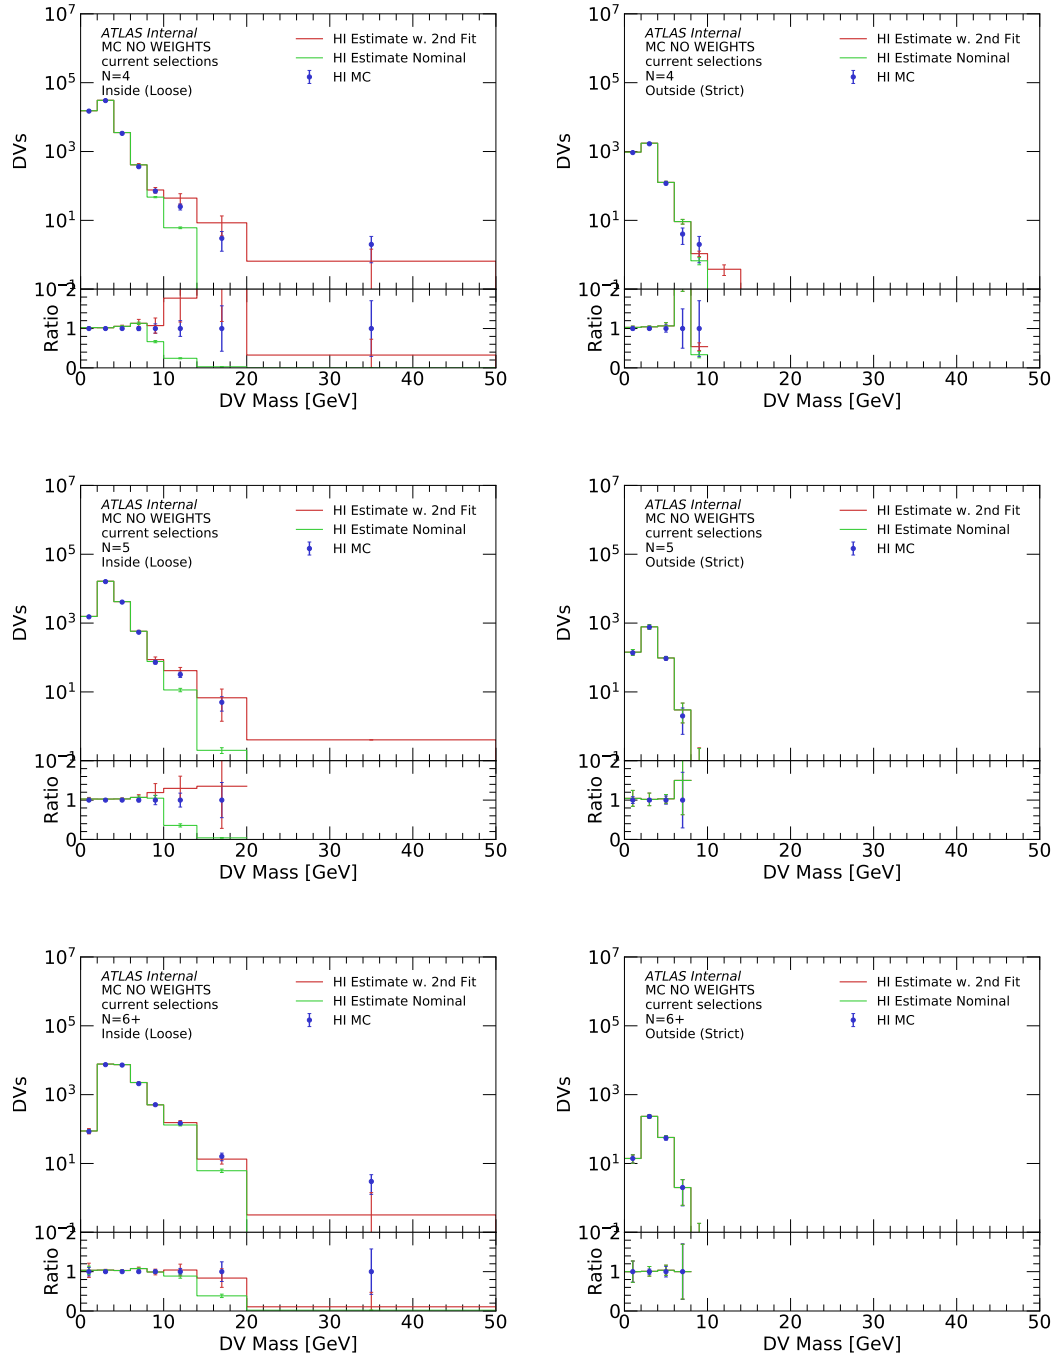


Figure 3.21: The estimate in the full MC against the tSP MC as a sort of closure test. The top row is for inside material and the bottom row is for outside material, with four (left), five (middle) or six or more tracks (right). Note that these do not have event weights applied and bin contents are not normalised to bin width. "Nominal" refers to the fit without the secondary component correction.




Figure 3.22: Visual representation of the procedure for estimating accidental crossings.

The procedure first involves obtaining a library of crossing tracks that can be added to DVs. This is done by taking $n_{\text{Tracks}}^{\text{DV}} = 3$ DVs where the invariant mass of 2 of the tracks lies within 0.05 GeV of the K_s mass of 0.497 11 GeV. This means that it is likely that this 3-trk DVs was actually a K_s decay that was crossed by an additional track. The fiducial volume of the detector is divided into 3 radial regions and 6 z regions during the entire procedure in order to account for the possible different behaviours of crossing tracks in the different regions. The momentum information of the track is saved as well as its information relative to its DV.

Therefore for each DV in a given region, a random track from the corresponding library is added to the DV, and the DV's properties are recomputed. The obtained mass distribution from this method is then normalised by the ratio $\frac{A_{K_s, 3\text{-choose-2 trks}}}{A_{K_s, 3\text{-choose-2 trks}} + A_{K_s, 2\text{ trks}}}$, where A_{K_s} represents the area under the K_s mass peak visible in the the given region. This represents the fraction of DVs that had a crossing track over those that could have a crossing track.

Due to discrepancies observed in the mass distribution between the true accidental crossing distribution in MC and the accidental crossings estimate in MC, as can be seen in Figure 3.25a and b respectively, the author investigated potential causes in order to resolve the difference. It was discovered by the author that the method by which a track is added to a DV is not trivial. Tracks have a momentum (p_x, p_y, p_z) , which will be added to the momenta of the other tracks in order to obtain the total invariant mass of the DV. A more convenient basis geometrically is to define them by their 3-momentum magnitude p , the component of that momentum transverse to the beampipe p_T , and the angle around the beampipe ϕ_{trk} .

The existing method determined the momentum of the track to add by keeping p and p_T consistent, and then changing only ϕ_{trk} such that the $\Delta\phi = \phi_{\text{trk}} - \phi_{\text{PVDV}}$

between that of the track and the PV-DV vector of the DV being considered is the same with the new DV, with respect to the DV the track originally came from.

As is visible in Figure 3.25b, the shape of the mass distribution obtained by this method in MC does not match the true distribution of accidental crossings in MC. This is because the mass of the DV, obtained by the norm of the sum of track momenta, is dependant on the direct dot product angles between the individual tracks, and is agnostic to any particular orientation. As seen in Figure 3.23, the projection angle ϕ is not necessarily related to the actual angle between the track and the other tracks in the DV or the PV-DV vector. This results in the tracks essentially being added at angles that are somewhat random, resulting in a smeared mass distribution with a thick tail and loss of recovery of behaviours of individual tAX components.

This could be fixed if instead the direct dot-product angle $\Delta\alpha$ between the track and the PV-DV vector was kept constant. However this is not trivial to do. The track p_T is an important parameter to attempt to keep constant since the momentum of the track is obtained via measuring the radius of curvature of the track caused by a charged particle in the transverse direction where the bending is caused by a magnet. This means that the uncertainties and cuts on tracks are all in fact based on the transverse momentum p_T , rather than the total scalar momentum. However since the total scalar momentum is what contributes to the DV mass, it would be ideal to keep both of these parameters constant.




Figure 3.23: Visual representation of a DV and track in the detector, showing the PV-DV vector and the angles $\Delta\alpha$ and $\Delta\phi$.

The constraint to keep a fixed total scalar momentum and a fixed p_T corresponds to a cone of possible track momenta around the collision axis. The constraint to keep a fixed $\Delta\alpha$ between the track and the PV-DV vector corresponds to a cone of possible

track momenta around the PV-DV vector. The intersection of these cones are our possible track momenta that we can set such that they will be properly added to the DV, as shown in Figure 3.24. The solution of this is as follows, where one only needs to determine ϕ as p_T and p are fixed and a modulus of 2π is implied:

$$\phi_{trk,new} = \phi_{trk,old} \pm \cos^{-1} \left[\frac{p_{trk} r_{PVDV} \cos(\Delta\alpha) - p_{trk,z} r_{PVDV,z}}{p_{trk,T} r_{PVDV,T}} \right]. \quad (3.9)$$

There are typically two solutions, the sign is chosen randomly. If there is no solution, a new track is drawn from the library. \vec{r}_{PVDV} represents the PV-DV vector.






Figure 3.24: Visual representation of the track momentum constraints. The blue cone represents the constraint from p and p_T . The green cone represents the constraints from p and $\delta\alpha$. The red line represents the PV-DV vector. The two intersections of the cones, in yellow, are the solutions for momenta satisfying all constraints.

After the implementation of this new method, Figure 3.25c shows a much better agreement with the shape of tAX, where the individual behaviours of the individual components can now be recovered.


The statistical uncertainties from the vertices the tracks are added to, are propagated to the final template. Uncertainties from the crossing-factor are determined by comparing them to when the crossing-factor is determined from $\Lambda \rightarrow p\pi$ decays instead.



(a) tAX derived directly from truth information (in other words, what we want to estimate)



(b) eAX as estimated by the old $\Delta\phi$ method



(c) eAX as estimated by the new $\Delta\alpha$ method

Figure 3.25: Mass distributions in MC divided by subcategory for AX, with what is aimed to be estimated and the two angular estimate methods. These plots are made without event weights.

3.4.3 Estimating Merged Vertices (eMV)

A method for estimating the merged vertex background predates the author's involvement in the analysis. Since merging of vertices occurs in the vertexing algorithm only if the distance significance squared S^2 between the vertices to merge is below 100 (unitless), this means that one can merge vertices from the same event and from different events and compare the distance significance distribution. One should observe a deficit in the same-event merged vertices in the $S^2 < 100$ region, as can be seen in Figure 3.26, corresponding to vertices that have already merged that we wish to study. A merged vertex template can be formed by merging vertices from different events, and then normalising this template to the distance significance deficit in the ratio of the same-event to different-event distance-significance distributions.




Figure 3.26: Comparison of the S^2 distribution from 2 and 3-track vertices that were merged together from the same vs. different events for the trackless SR. The merging portion of the secondary-vertexing algorithm merges vertices with $S^2 < 100$. The deficit in the $S^2 < 100$ distribution for the same-event merging represents vertices that were already incorrectly merged by the algorithm. This plot was produced by Atsushi Mizukami and Emily Anne Thompson, as it appears in the analysis paper [1]

3.4.4 Closure Uncertainty for Accidental Crossings (eAX) and Merged Vertices (eMV)

Accidental crossings and merged vertices cannot be separated in truth information due to the lack of information about pileup and the internal workings of the vertexing algorithm. This is because the vertexing algorithm is run during the reconstruction of the data files which is done centrally within the collaboration, and this information is not saved to the file by the algorithm. For example, DVs that are SP with two extra pileup tracks cannot be categorized with certainty as they can be either category, depending on the physics of the pileup tracks not available to us. It could be that if the two pileup tracks represents an SP pileup vertex that this is a true merged vertex. However if they are separate tracks in the pileup event this would be a double accidental crossing. Additionally selections have an effect. For example, if one has a merged vertex, however only one track of one of the vertices merged survives selection, it would not be possible to tell whether or not such an event should be categorized as a merged vertex or accidental crossings.

Therefore, for closure tests of the estimation methods, ideally this would be done separately for SP and other DVs. The closure for SP has been detailed in Section 3.4.1. A separate closure test is then done for the combined eAX and eMV estimates, comparing them to all non tSP DVs in the truth information, which can be seen in Figure 3.27. Since the SP DVs dominate in low mass to the point that a non-closure uncertainty on the other components would not matter, a flat uncertainty for the other components is derived by only comparing the high mass bins where the eAX and eMV dominate. There is an underestimation, and so a flat uncertainty of 40% corresponding to this underestimation is assigned to the combined eAX and eMV in the final estimate. This is derived from the inside material regions as an accurate non-closure cannot be derived in the high-mass outside material regions due to large statistical uncertainties.

The non-closure is less than with the original ϕ -angle method, where the ϕ -angle method would require a non-closure uncertainty of 70%. However, due to severe time constraints in the analysis the official result from the paper uses the original ϕ -angle method. This was justified due to the fact that the ϕ -angle method is more conservative in the SR (due to DVs that would have otherwise been in the low mass region peak being shifted to the high mass region) and it still remaining within a suitable uncertainty. This is intended to be fixed in the future DV+MET analysis,

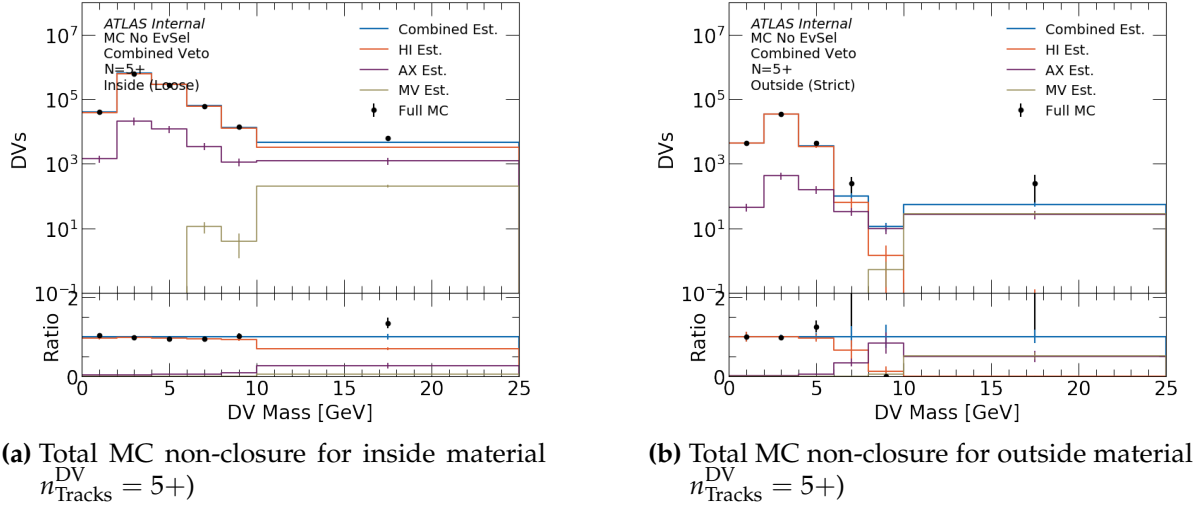


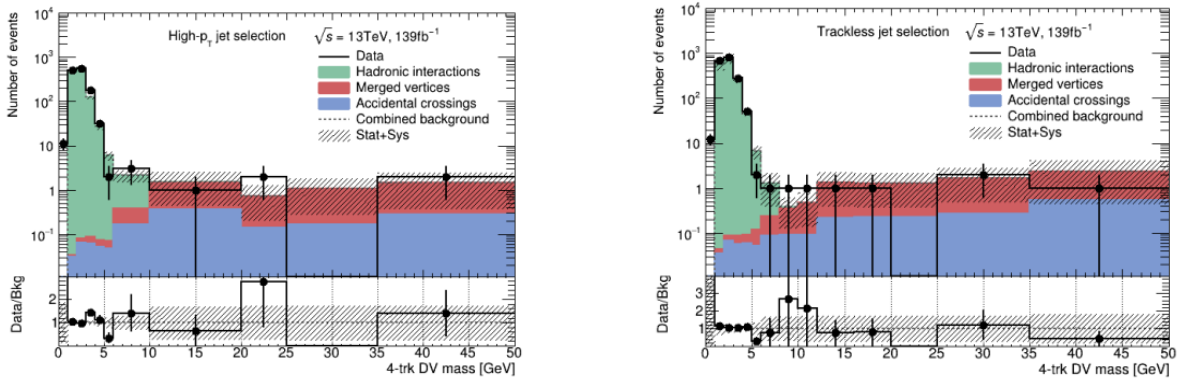
Figure 3.27: The combined estimate performed in MC is compared to true MC. The difference in the high mass bins is taken to be a non-closure uncertainty for the combined eAX and eMV. The overall systematic for the eAX trace was developed by Emily Anne Thompson. The trace for the eMV estimate is produced by Emily Anne Thompson and Atsushi Mizukami

however for the rest of this chapter, results will be stated using the ϕ -angle method in order to be consistent with the official paper results.

3.4.5 Combining Background Estimates

Several methods of combining the separate background estimates were investigated, primarily those that involved allowing the normalisations of the respective estimates to be determined via a simultaneous fit of all of the estimate shapes in data. However it was determined by the author that the methods which involved simultaneous fits were not reliable and yielded unstable results. This was due to the transition period between where hadronic interactions dominate towards where the other backgrounds dominate being quite close to the $m_{\text{DV}} = 10 \text{ GeV}$ cut-off, meaning there was insufficient information about the behaviour of the accidental crossings and merged vertices shapes to properly fit their normalisations. As a result, the final method chosen has only the hadronic interactions estimate directly fitted in data, primarily due to it being fairly challenging to derive a normalisation otherwise. The other estimates use their own normalisations as described in their respective sections.

The combined background estimate is initially validated during the development of the estimate in the unblinded $n_{\text{Tracks}}^{\text{DV}} = 4$, $m_{\text{DV}} \in [10 - 20 \text{ GeV}]$ region. The estimate is finally validated, right before the unblinding of the SR, in the blinded $n_{\text{Tracks}}^{\text{DV}} = 4$, $m_{\text{DV}} > 20 \text{ GeV}$ validation region. The results are seen in Figure 3.28, created by the team using inputs from those responsible for each individual estimate, taken from the paper [1], where good agreement is seen.



(a) High p_T , outside material, $n_{\text{Tracks}}^{\text{DV}} = 4$ region, where the validation region is $m_{\text{DV}} > 20 \text{ GeV}$.

(b) Trackless, outside material, $n_{\text{Tracks}}^{\text{DV}} = 4$ region, where the validation region is $m_{\text{DV}} > 20 \text{ GeV}$.

Figure 3.28: The blinded validation region in data ($n_{\text{Tracks}}^{\text{DV}} = 4$, $m_{\text{DV}} > 20 \text{ GeV}$), against the combined background estimate, created by the team and taken from the official paper [1].

The final estimated backgrounds are 1.08 ± 0.69 for the High p_T SR and 2.1 ± 1.1 for the Trackless SR, as calculated by the team from respective inputs in the official paper [1].

3.5 Inclusive Background Estimates

The combined background estimate is complex and with sizable uncertainties. Another estimation method was developed by the analysis team in order to estimate background DVs inclusively, agnostic to their particular type. As this analysis has quite unique backgrounds, these separate estimates are then cross-checked against each other in order to instil confidence in the understanding of the estimate.

This inclusive background estimate method is based on all of the background types being correlated with the presence of tracks, and therefore track-jets. In other words,

the more track-jets there are in the events, the more tracks there are, and therefore the more DVs. Therefore, if one can determine how many background DVs appear per jet in some control region, knowing the number of jets in the signal region one can therefore determine the number of expected background DVs. The control region is chosen to be single-photon-triggered events that fail both the High p_T and Trackless jet requirements.

These track-jets are reconstructed slightly differently than typical jets reconstructed from calorimeter deposits. These are reconstructed with an anti- k_t algorithm with $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} = 0.4$ from all $p_T > 1$ GeV and $|d_0| < 2$ mm tracks. The d_0 requirements makes the track-jet insensitive to LLP decays.

In the control trigger region, 2D binned histograms are parameterized as a function of the closest track jet p_T (p_T^{jet}) and the number of tracks in the closest track jet ($n_{\text{Tracks}}^{\text{jet}}$), are made for the number of track jets matched to a DV and for the total number of track jets. Their ratio, which represents the probability of a DV being present given a jet is present, is referred to as the jet-DV probability,

$$P(DV|jet(p_T^{\text{jet}}, n_{\text{Tracks}}^{\text{jet}})) = \frac{\# \text{jet}(p_T^{\text{jet}}, n_{\text{Tracks}}^{\text{jet}}) \text{ matched to DVs}}{\# \text{jet}(p_T^{\text{jet}}, n_{\text{Tracks}}^{\text{jet}}) \text{ total}}. \quad (3.10)$$

In order to reduce limitations due to a small number of events, the control trigger region uses the region $n_{\text{Tracks}}^{\text{DV}} \geq 4, m_{\text{DV}} > 5$ GeV for its DVs. As the signal region may be defined in an $(n_{\text{Tracks}}^{\text{DV}}, m_{\text{DV}})$ region that is not the $(n_{\text{Tracks}}^{\text{DV}} \geq 4, m_{\text{DV}} > 5 \text{ GeV})$ used in the control trigger region, an f factor is calculated to compensate for this. This normalises the result by the ratio of the number of events with DVs falling into the $(n_{\text{Tracks}}^{\text{DV}}, m_{\text{DV}})$ region in the control trigger region to the number of events with DVs falling into the $(n_{\text{Tracks}}^{\text{DV}} \geq 4, m_{\text{DV}} > 5 \text{ GeV})$ region in the control trigger region.

Therefore, the estimate method can be summarised as follows:

$$N_{SR, (n_{\text{Tracks}}^{\text{DV}}, m_{\text{DV}})} = f_{(n_{\text{Tracks}}^{\text{DV}}, m_{\text{DV}})} \cdot \sum_{\text{jet}_{SR}(p_T^{\text{jet}}, n_{\text{Tracks}}^{\text{jet}})} P(DV|jet_{SR}(p_T^{\text{jet}}, n_{\text{Tracks}}^{\text{jet}})) \quad (3.11)$$

$$f_{(n_{\text{Tracks}}^{\text{DV}}, m_{\text{DV}})} = \frac{N_{CR, (n_{\text{Tracks}}^{\text{DV}}, m_{\text{DV}})}}{N_{CR, (n_{\text{Tracks}}^{\text{DV}} \geq 4, m_{\text{DV}} > 5 \text{ GeV})}}. \quad (3.12)$$

A closure test with MC in the SR saw good agreement. This method was validated in both MC and data in various $(n_{\text{Tracks}}^{\text{DV}}, m_{\text{DV}})$ regions outside of material for both multi-jet selections, particularly with $(n_{\text{Tracks}}^{\text{DV}} = 4, m_{\text{DV}} \in [10, 20]\text{GeV})$ and $(n_{\text{Tracks}}^{\text{DV}} \geq 5, m_{\text{DV}} \in [5, 10]\text{GeV})$. This saw good agreement except for a consistent overestimation of background in the $(n_{\text{Tracks}}^{\text{DV}} \geq 5, m_{\text{DV}} \in [5, 10]\text{GeV})$ regions but only in MC. This region is dominated by hadronic interactions, so this hints that this method may not perform very well in hadronic interaction rich regions. An uncertainty is derived from this non-closure.

Systematic uncertainties in the control region, the possibility of a dependence on the number of track-jets in the event, and pileup were also taken into account.

The final estimated yields for background events in the signal region are $0.46_{-0.30}^{+0.27}$ for the High p_{T} SR and $0.83_{-0.53}^{+0.51}$ for the Trackless SR, as determined by Emily Anne Thompson and Moe Wakida in the paper [1].

3.6 Uncertainties

In addition to the uncertainties described by the background estimation methods, since the results are derived from comparing MC signal yields to data, uncertainties relating to difference in data and MC are important to take into account. Effects of pileup were determined to be negligible enough to be ignored. Uncertainties relating to the tracking uncertainty, jet energy scale and resolution, data integrated luminosity, and MC modelling of initial state radiation were considered. The data integrated luminosity uncertainty is 1.7% as determined by the collaboration. The jet energy scale and resolution uncertainties as a result of the jets originating from displaced decays were found to be negligible in the Strong RPV signal model, but up to 10% in the EWK RPV signal model. This is due to the EWK RPV signal model primarily relying on the Trackless Jet SR, whose required trackless jets are the focus of this uncertainty. The uncertainty from initial state radiation was determined by varying parameters in the generator and showering during production and comparing to the nominal. These were determined to be negligible for the Strong RPV signal model and up to about 20% for the EWK RPV signal model. These jet uncertainties were determined by Giulia Ripellino in the paper [1].

The uncertainty on standard tracking reconstruction efficiency has already been determined by the collaboration to be approximately 1.7% [69]. Therefore, by plotting the radial distribution of K_s decays in both data and MC, normalising in the low radius region where the efficiency is well understood, one can estimate the MC mismodelling of LRT tracking. This is added in quadrature to the uncertainty from standard tracking to obtain a per-track efficiency uncertainty, which can be propagated to the final event counts. This was determined to be between 14% and 17% by Giulia Ripellino in the paper [1].

3.7 Results

The predictions of the two estimation methods agree with each other, lending confidence that the estimates are reliable and correct.

After unblinding, in data 1 event was observed in the High p_T SR and 0 events were observed in the Trackless SR. The estimates and final data yields are summarised in Table 3.9. This one observed event is from 2017 data, and represents a DV with 5 tracks and $m_{DV} = 32.6$ GeV, at $R_{xy} = 7.5$ mm, which lies inside the beampipe. This event is shown in the event display in Figure 3.29.

Signal Region	Combined	Inclusive	Observed
High p_T	1.08 ± 0.69	$0.46^{+0.27}_{-0.30}$	1
Trackless	2.1 ± 1.1	$0.83^{+0.51}_{-0.53}$	0

Table 3.9: Summary of the yields in the signal regions as estimated by the combined and inclusive estimates, compared to the final unblinded yields in data. These are the final numbers from the paper, calculated by various members of the analysis team [1].

As the observed data is consistent with the background estimate, and therefore the Standard Model, exclusion limits on the BSM scenarios are set. Limits are calculated according to the CL_s method [70] using pyHF [71][72]. In general, the Strong RPV model has its sensitivity primarily in the High p_T region, whereas the EWK RPV model has important contributions from both. As the two SRs are orthogonal, limits can be combined as a two-bin fit. The final exclusion limits of long-lived particle mass vs.




Figure 3.29: Event display of the single observed event in the signal regions in data. Created by Moe Wakida. The blue circles represent different proton-proton interactions along the beamline. The green tracks are the reconstructed DV. The shaded brown cylinder represents the beampipe.

lifetime on both models are shown below in Figure 3.30, where the Strong RPV model is plotted for a fixed gluino mass of $m_{\tilde{g}} = 2400$ GeV.




Figure 3.30: The dashed line and the yellow band represents the expected limits and its $\pm 1\sigma$ uncertainty. The solid red line represents the observed limit, with the dotted red lines representing the theoretical $\pm 1\sigma$ uncertainty. These plots are taken directly from the paper [1], created by the team using respective inputs, using code created by both Stefan Richter and myself.

The observed limit appropriately does not deviate largely from the expected limit. For the Strong RPV model, this limit therefore excludes the possibility of nearly the entire lifetime range from 0.01 to 1 ns of long-lived neutralino mass down to a couple hundred GeV. For the EWK RPV model, this excludes up to at least 1.2 TeV of long-lived neutralino mass for the lifetime range from 0.01 to 10 ns.

This is the first ATLAS search result for such multijet-triggered high mass and high number of track displaced vertices with $\sqrt{s} = 13$ TeV data, as well as the first ATLAS search to have set limits in such long-lived RPV scenarios. The CMS experiment had recently conducted a similar search, with a model similar to our EWK RPV model, however the primary difference is that their fiducial volume was limited to a radius of less than 20.9 mm such that they do not need to consider a material map. At long-lived neutralino lifetime of $\tau = 0.03$ ns, their analysis was able to exclude up to 1.3 TeV as opposed to this analysis which was able to exclude up to 1.5 TeV [73].

This analysis in theory should have sensitivity to Higgs Portal models where the scalars decay into $b\bar{b}$ pairs at the displaced vertices, which of course can also count as jets themselves. Many of the Higgs production modes can also have additional jets from the production process itself. The $t\bar{t}H$ mode in particular will have jets due to the b 's and possibly from the W 's of the t quark decays. The VBF mode has a high

cross-section as well as forward jets from the production process. The ggH mode, although to within leading order has no jets from the production process, may have initial state radiation and has the dominating cross-section.

As the jet requirements used in this analysis are so stringent, very few MC events produced would pass the analysis requirements, meaning that a dedicated filter has to be produced for these samples and a very large number of events need to be available as input to the filter. Unfortunately due to technical limitations, these samples were unable to be completed within the timespan of this thesis. However, from preliminary numbers with a reduced MC sample test obtained in the formulation of the sample request by Anna Mullin and myself, we expect a yield in the Trackless SR of 59.6 ± 21.8 for one of the most sensitive points of $m_S = 45 \text{ GeV}$ and $c\tau = 10 \text{ mm}$ when we include all Higgs production modes at 100% Higgs branching ratio. Note that the large statistical error in the MC are due to the limit number of raw events available in this small test sample. A limited plot of the expected yields for various signal points in the Trackless SR at 100% Higgs branching ratio inclusive of all Higgs production modes is shown in Figure 3.31. Effectively no sensitivity is expected from the High- p_T SR.




Figure 3.31: Expected yields for the Higgs Portal model at 100% Higgs branching ratio inclusive of all Higgs production modes for a limited signal scan of LLP scalar mass and lifetime. Samples used to create this plot were produced with Anna Mullin.

As a rough estimate one could consider the uncertainties of the background estimate of the analysis to be able to reject a signal of 3 DVs to a 95% certainty. Taking as a rough estimate the nominal value of the signal yield without its uncertainties, this would allow this analysis to roughly exclude down to the scalar occupying 5% of the Higgs branching ratio for the $m_S = 45 \text{ GeV}$ and $c\tau = 10 \text{ mm}$ point. The current limits on this point for hadronically-decaying scalars is also approximately around 5% from independent ATLAS [74] and CMS [75] analyses, meaning there is the potential for a statistical combination with a reinterpretation of this analysis to obtain the best limits to date on this branching ratio.

Although there is sensitivity to Higgs Portal models in this analysis, which triggers partially on jets from the Higgs production process, the jet requirements need to be very stringent which limits the sensitivity. In the next chapter a new analysis is introduced, the DV+MET search, where one triggers instead on MET from the event. Here, DVs will still be produced by the decay of the long-lived dark sector particles that the Higgs decays into. However, as many Higgs production processes also produce neutrinos, and the MET requirements are much less stringent than those on jets, there is the potential to have even more sensitivity with this alternative approach.

Chapter 4

The Displaced Vertex + Missing Transverse Energy Search (DV+MET)

“What we lack in experience we make up for in enthusiasm.”

— Fionn Bioshop, PhD Particle Physics, LHCb, University of Cambridge

This chapter details the "DV+MET" search, where new long-lived particles (LLPs) Beyond the Standard Model are searched for using displaced vertex signatures in events that are triggered by missing transverse energy (MET). MET is simply used as a trigger and there are no additional requirements that it is directly originating from the LLP decay DV or not. Similar to the DV+Jets analysis, this is to allow the analysis to be as general as possible, and inclusive to a wide variety of BSM models. The primary physics motivation for using this specific trigger was simply because the Run 2 analyses triggering on muons [57] and jets [1] have already been completed, leaving MET as the next most motivated general trigger. This search uses the full Run 2 proton-proton collision data set of 139 fb^{-1} at centre-of-mass energy $\sqrt{s} = 13 \text{ TeV}$. This is the first full Run 2 DV+MET search, following up on the early partial Run 2 DV+MET analysis [2].

In this updated search, in addition to using what is now the standard VSI vertexing algorithm (described in Section 2.3.4), a new fuzzy vertexing algorithm developed by Risa Ushioda is being used for the first time in ATLAS. If an LLP decays into heavy quarks, such as b -quarks, the heavy quarks have a lifetime that results in their hadrons decaying a few mm after they are produced in the DV. This means that the stable

tracks that are reconstructed by the detector will not all point back to the original DV, but rather nearby. The traditional VSI algorithm requires that all tracks point back to the DV. This requirement may be too strict for these heavy decay DV cases. The fuzzy vertexing algorithm aims to solve this by loosening the restriction that tracks must originate from the DV, instead allowing them to come from a volume around the DV. This is illustrated visually in Figure 4.1.




Figure 4.1: An LLP originating from the PV decays into a $b\bar{b}$ pair, which each decay later into several tracks. Black dotted tracks are not reconstructed by the detector. Orange dotted tracks are not reconstructed into the DV, but orange solid tracks are. The blue circle represents the DV, which is point-like in VSI vertexing and a volume in fuzzy vertexing.

Unfortunately, the loosening of the vertexing requirements also results in increasing backgrounds. To ensure no loss of sensitivity to models with LLPs decaying into light quarks, three different signal regions are planned to target different areas of phase space, searching for events with 1) at least 1+ VSI DV (the "1+ VSI DV SR"), 2) exactly 1 fuzzy DV (the "1 fuzzy DV SR"), and 3) 2+ fuzzy DVs (the "2+ fuzzy DV SR"). Although the analysis is signature-driven, these three SRs are each optimized using a different particular signal benchmark model.

The 1+ VSI DV SR uses an R-Hadron R-parity-conserving (RPC) SUSY model. This model is used as it is the same model used in the early partial Run 2 analysis [2], therefore allowing for a direct comparison in sensitivity gain. Here a gluino pair is produced from the collision. These are hadronized into R-hadrons (hadrons with a SUSY particle in them), hence the name of the model. These then decay into squarks. The squark then decays into a neutralino and quarks. The squark is assumed to be

significantly heavier than the gluino, meaning it is produced very off-shell, causing the gluino to become long-lived. This decay of the long-lived R-hadron gluino into the neutralino and quarks is what creates the displaced vertex signature. The neutralino is invisible and constitutes the MET that is triggered on [76]. The Feynman diagram for this process is shown in Figure 4.2a.

The 1 fuzzy DV SR uses a Wino-Bino RPC SUSY model as it has a single heavy-flavour DV in its final state. Here a $\tilde{\chi}_2^0$ is produced in the collision and decays into a $\tilde{\chi}_1^0$ and a very off-shell Higgs. The $\tilde{\chi}_1^0$ is stable and invisible, and comprises the MET the event is triggered on. The long-lifetime of the $\tilde{\chi}_2^0$ originates from it only being able to decay via a very off-shell mediator, the Higgs. Additionally the mass splitting between the neutralinos is assumed to be on the order of 30 GeV which further prolongs its lifetime. The off-shell Higgs decays into a $b\bar{b}$ pair, forming a single fuzzy DV [77]. The Feynman diagram for this process is shown in Figure 4.2b.

The 2+ fuzzy DV SR uses a Higgs Portal model as there will be two heavy flavour DVs in its final state. A Higgs boson is produced in the collision via any of the Higgs production processes. Neutrinos from the Higgs production process, particularly from a Z or W decay in the VH or ttH modes, will be the source of MET that can be triggered on. The Higgs decays into a scalar LLP pair, which then each decay into a $b\bar{b}$ pair. Each of these $b\bar{b}$ pairs will form a displaced fuzzy vertex, assuming that they decay within the fiducial volume. No explicit mechanism for the long-lifetime of the scalar is imposed. The Feynman diagram for this process is shown in Figure 4.2c. One may note that if one were to design an analysis specifically to target this model, one would rather trigger on a single lepton as any of the W decays that produce a neutrino will also produce a lepton, and the lepton triggers are more reliable. However this case is already covered by the already completed DV+muon search [57] and the already currently underway Higgs Decay with Multiple Implementations analysis [78]. This analysis serves primarily to help complete the phase space of possible general trigger signatures in a model agnostic way, and this model is used here as a benchmark rather than as motivation.

I was analysis contact for this analysis. I therefore lead the analysis team together with my co-analysis contact Evan Michael Carlson. A large part of my responsibility on the analysis is coordinating the team, designing the analysis, and providing feedback to students. Another large part of my responsibility is closely providing support and supervision to other students on analysis to help them complete their tasks.

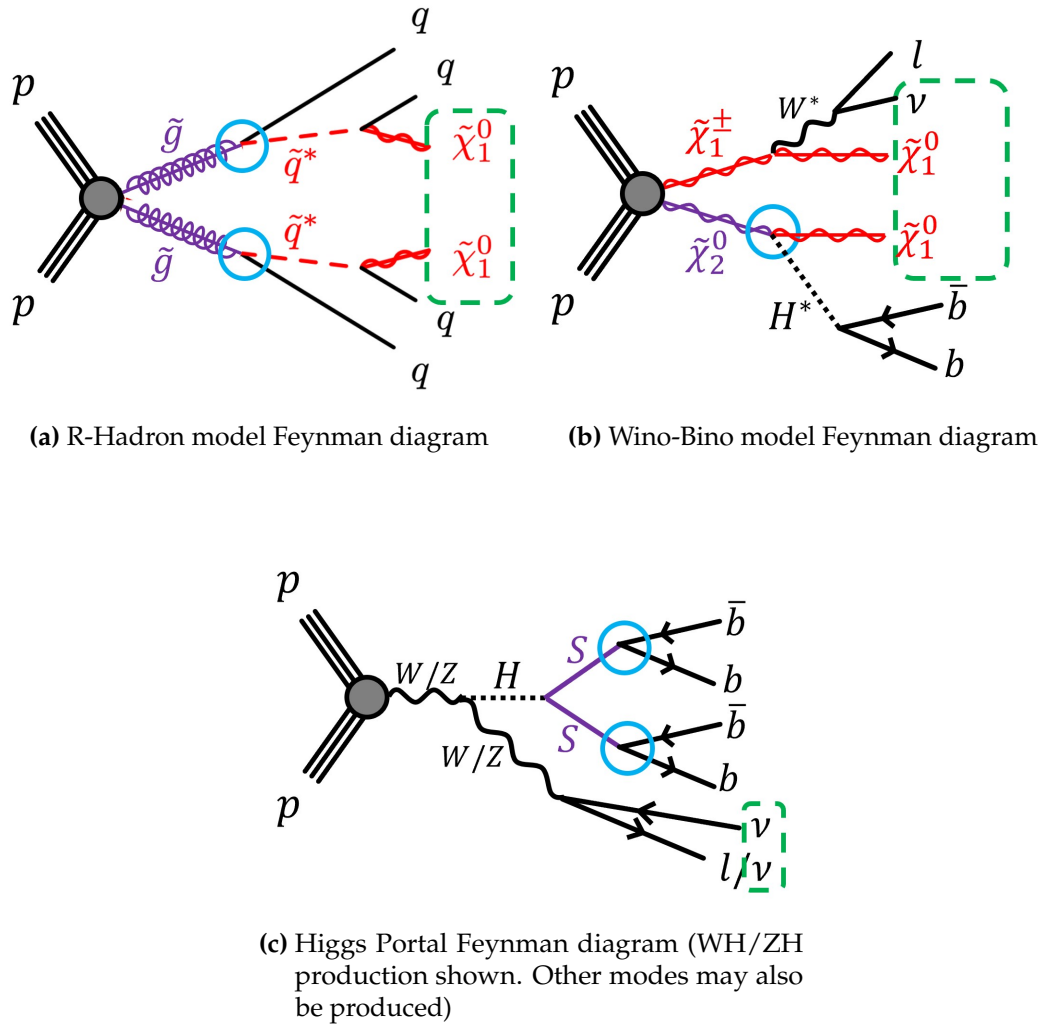


Figure 4.2: Benchmark models for the DV+MET analysis. The blue circles represent the DVs. The object surrounded by the green dots represent objects that the analysis triggers on, in this case MET.

My primary tangible contribution to the analysis were in updating the material map veto system, doing track cleanings and signal region definitions for the 2+ fuzzy DV SR, and together with Anna Mullin developing the MC request for the Higgs Portal model. I am also heavily involved in the development of an updated inclusive background estimate in collaboration with others in the analysis team. I also did preliminary track cleaning studies for the 1+ VSI DV SR before this was taken over by another person.

As my primary contribution was on the 2+ fuzzy DV SR, this is what this chapter will focus on. The analysis is still in progress at the time of this writing, so only current progress on the material map and sensitivity studies for the 2+ fuzzy DV SR will be covered.

4.1 Samples and Basic Selections

The full Run 2 dataset of 139 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13 \text{ TeV}$ was used for this analysis. The data directly analysed in this analysis is collected by a set of MET triggers. However, data collected from other triggers will also be used in the analysis for background estimation. Like the DV+Jets analysis, the standard processing for the particular software release used at the time of this analysis ("Release 21") did not include the large-radius tracking required to reconstruct displaced vertices (as described in Section 2.3.4). The same filter derived by the RPV/LL group to select events in data that would be interesting for several different analyses to undergo bulk reconstruction.

The various MET triggers themselves vary depending on physics data-taking period, and typically trigger on a MET between 90 GeV to 120 GeV. This online MET is reconstructed during data-taking and so is done with a fast but less accurate algorithm. The MET reconstructed for use in the analysis is done offline, after data-taking with a more accurate algorithm and taking into account calibrations. Typically this means that there will be a turn-on curve in the offline MET, as some events that lie at the cut-off of the online MET may in fact be misreconstructed and end up on either side of the cut-off in the offline MET. Due to this, the early Run 2 analysis placed their offline MET cut-off at 250 GeV to ensure they were at the full efficiency of the trigger [2]. In this current analysis we are investigating the use of this trigger turn-on region and lowering this threshold. However, there is a limit due to selections already made

on events in order to pass the RPV/LL filter before they can be re-reconstructed with LRT. These limits impose that $\text{MET} > 150 \text{ GeV}$ and an internal intermediary rough form of offline MET referred to as $\text{MET_LHT} > 180 \text{ GeV}$. MET_LHT ("LocHadTopo", meaning local hadronic topological cluster calibration) is a more rougher version of MET, but calculated only using calorimeter cells, where the energy is calibrated by first classifying clusters as being hadronic or electromagnetic [79]. These limits are what have been defined as tentative event selections as of this writing. Naturally, the event must also contain DV(s) that pass the selections of either the 1+ VSI DV, 1 fuzzy DV, or 2+ fuzzy DV SRs. These selections are outlined in Table 4.1. The VSI and fuzzy SRs are designed to focus their sensitivity to light and heavy decay models respectively, and are expected to not have significant cross-sensitivity. For this reason, the fuzzy and VSI SRs are not kept orthogonal.

Event Selections:

- Passes MET triggers
- $\text{MET} > 150 \text{ GeV}$
- $\text{MET_LHT} > 180 \text{ GeV}$
- Contains DV(s) that passes the DV selections of either the 1+ VSI DV, 1 fuzzy DV, or 2+ fuzzy DV SR.

Table 4.1: DV+MET Event Selections

Simulated events for the signal models are needed in order to estimate yields, needed to both design signal regions and draw conclusions about the results. The Higgs Portal model was generated by taking MC of Higgs boson generation already centrally produced in ATLAS, and then manually using Pythia 8.307 to decay the Higgs boson into 2 long-lived scalars, those long-lived scalars into a $b\bar{b}$ pair each, and then to perform the showering and hadronization of the event [64].

In theory, any Higgs production mode has the possibility of creating sufficient MET to pass the trigger if the LLP decays past the calorimeter or if the LLP decay products are displaced such that they are counted as MET. However, only the $t\bar{t}H$ and VH productions are focused on in this work as these are the only Higgs production modes that create neutrinos. The VH production modes were produced using POWHEG 2 and Pythia 8.212 at next-to-leading order [80][64]. The $t\bar{t}H$ production modes were produced at leading order using POWHEG 2 [80].

Much like the DV+Jets analysis, although the background estimate is largely data-driven, background simulated events are still needed to study and understand the backgrounds, as well as to test if the background estimates function properly. The dominant background for a MET-triggered analysis will be $Z \rightarrow \nu\nu$ events. Therefore, $Z \rightarrow \nu\nu$ events were generated using Sherpa 2.2.1 at leading order [81]. $W \rightarrow l\nu$ also will exist as a background, however it is expected to not be as high as leptons are included in the MET reconstruction used in the analysis. Those samples were not produced due to the high computational cost of running MC samples through LRT and secondary vertexing requiring us to prioritize samples.

The effects of pileup is generated in the same manner as with DV+Jets. Pileup interactions are generated with Pythia 8.210 [64] and overlaid on top of the samples in three different profiles that correspond to different data-taking periods: MC16A (2015-2016), MC16D (2017), and MC16E (2018). Including all relevant Higgs production modes, 180,000 events were produced for MC16A, 260,000 events were produced for MC16D, and 300,000 events were produced for MC16E. For the background samples, approximately 2.9 million events were produced for MC16A, 3.2 million events were produced for MC16D, and 4.2 million events were produced for MC16E. All MC samples are propagated through a simulation of the ATLAS detector [65] using GEANT4 [66], after which the samples are processed in the same manner as data.

4.2 Updating the Material Map Veto

As discussed in Section 3.2.3 in the DV+Jets analysis, a material map veto is essential for the analysis as it gives a reliable way of removing what would otherwise be a dominating hadronic interactions background.

No veto is perfect, and there will always be regions of the fiducial region in which one cannot be certain whether or not lies inside detector material or not, particularly in literal edge cases where the decision becomes very sensitive to the material map accuracy, DV resolution, and binning. One could be conservative in developing the veto, deeming any uncertain region as inside material, in order to get a clean signal region as free of hadronic interactions as possible. However this also means that in the inverse veto region that there may be an increase in the possibility of signal contamination as the DVs in regions actually outside of material that are miscategorized as inside material will be dominated by signal. This is unideal as the inverse veto

region is needed for validation and control regions. The pragmatic, albeit somewhat inelegant, solution for the purposes of analysis is to develop two vetoes: a loose veto whose inverse will be used for "inside material" regions, and a strict veto which will be used to define "outside material" for the signal region.

Additionally, the veto needs to be different if it is applied in data or MC. There are two major differences to keep in mind: the GEANT4 map used in the simulation may be different in real life, and as mentioned in Section 3.2.3 the beampipe and pixel layers are not perfectly at the centre of ATLAS in real-life. The coordinate reference for ATLAS is measured from the cavern itself. In the fiducial volume there lies the pixel detectors which are attached to the beampipe which will therefore also be slightly offset in real life, and the SCT support attached to the ground which is stable.

The solution to the second problem is simple. The offsets of the beampipe and pixel layers are known and were determined by the collaboration after the alignment and track calibrations by observing what offset to the pixel layers was required to maintain their rotational symmetry. These offsets can be applied to the maps in the data vetoes. MC vetoes will not apply these offsets as the GEANT4 map used in simulation assumes perfect alignment.

As for the first problem, originally the solution was to use a data-derived map for data vetoes and to use a modified version of the simulation GEANT4 map for MC vetoes. The data-derived map and GEANT4 map are shown in Figures 4.3a and b respectively. In the GEANT4 case, different geometry regions would simply be tagged as either "material" or not material. As described in Section 3.2.3, for the data-derived map case in the outer region where there are insufficient DVs in data to derive a high resolution map, regions of the detector are blocked out using the GEANT4 map as guidance and the presence of DVs as confirmation [57].

In past analyses, the GEANT4 map would be directly used as the MC veto, and the data-derived map would be directly used as the loose data veto [57]. However, over the course of the DV+Jets analysis, two issues with this approach were discovered. The first is that there were clusters of DVs were discovered that were deemed "outside material" by the GEANT4 map, which corresponded to the locations of cooling-liquid tubes in the pixel detector. This is shown in Figure 4.4, where the inside of the tubes filled with liquid appears to the map as "outside material" at around $R_{xy} = 60$ mm. Additionally, the GEANT4 and data-derived maps appear visually different, vetoing a visually drastically different amount of the fiducial volume. This makes the data and

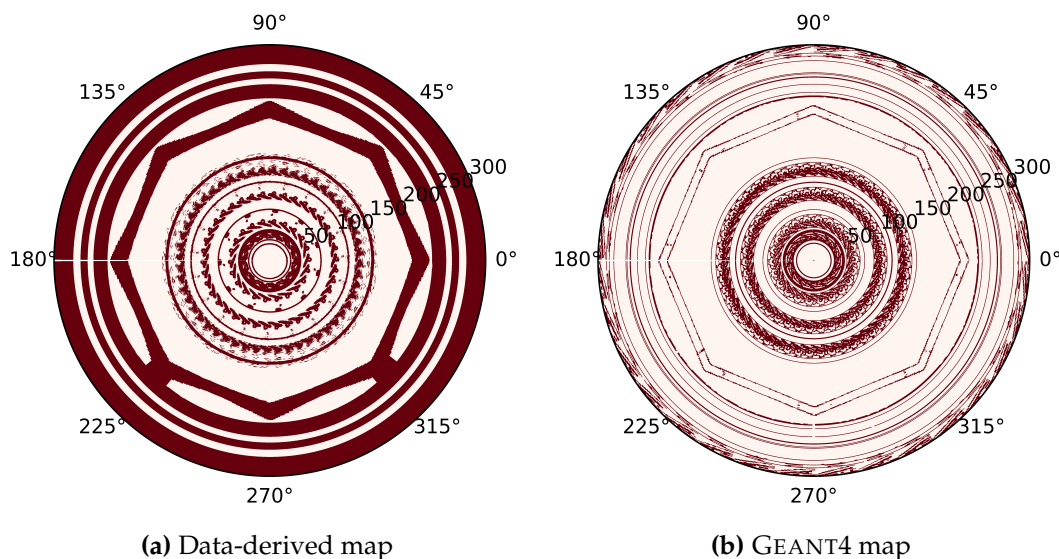


Figure 4.3: Visualizations of the data-derived and GEANT4 maps that are used for vetoes. These are shown in $(R_{xy}, \phi, z = 0 \text{ mm})$, with "inside material" denoted in red.

MC behave very differently. This is particularly detrimental when calculating signal yields and doing initial studies trying to develop tentative signal region definitions using MC background samples as a guide. We also observe regions in the MC map in the inner radial regions that are not in the data-derived map, as shown in Figure 4.5, which appear to be largely structural.

Ideally, the MC and data vetoes would be placed on equal footing. The GEANT4 map was designed in the first place to model the detector material regions as accurately as possible for simulations. Although this is not perfect, it is not expected to be significantly different from the detector at the resolution of the data-derived map. There is of course no way of directly validating the accuracy of the data-derived map in real-life, as the closest most accurate map we have of the detector is the GEANT4 map in the first place. Unfortunately, the timeline of the analysis, as well as computational constraints, do not permit the development of a "MC-derived" map to be made in the same procedure as the data-derived map. Additionally, getting access to the original GEANT4 map to remake the map is non-trivial. The solution agreed upon was therefore be to simply use the data-derived map for both MC and data. This means that the total fraction of the fiducial volume that is vetoed by the MC and data vetoes will be the same. However, of course the geometry actually being simulated in MC may be slightly different than the real-life geometry so the efficiency may be different.




Figure 4.4: A visualization of the GEANT4 map in the inner radial region with ($R_{xy}, \phi, z = 0$ mm). Yellow denotes "inside material". DVs with 5+ tracks from a dijets MC background sample are shown as white crosses overlaid on the map only if they happen to lie in bins the map deems as outside material.




Figure 4.5: Visualizations of regions in red that the GEANT4 map deems as "inside material" but the data-derived map deems as "outside material".

Therefore, four different vetoes will be developed: MC Loose, MC Strict, Data Loose, and Data Strict. The only differences between the data and MC vetoes will be that the pixel layer offsets will be applied for the data vetoes. The strict vetoes will start with the loose vetoes as a base, adding additional algorithms that can only further veto DVs that have not already been vetoed by the loose veto. The loose vetoes simply implement the data-derived map directly, vetoing DVs as "inside material" if their position lies in a map bin declared by the map as "inside material".

The strict vetoes will have three additional steps to further veto DVs. The first step, referred to as the GEANT4 map step, is to additionally veto DVs in regions the GEANT4 map deems as "Inside Material". This ensures that no regions of the fiducial volume which theoretically should have material are deemed as "outside material" due to the data-derived map statistically not being able to correctly identify this region as "inside material". This takes the risk that there may be no material there in real life in the case of any flaws in the GEANT4 map.

The second step, referred to as the neighbouring bins step, is to additionally veto a DV if any of the bins in the data-derived map surrounding the bin in which the DV lies is deemed as "inside material". If the edge of a material cuts through a bin but not enough of the volume of the bin lies inside the material, the entire bin may be deemed "outside material". However, this is not true as DVs that fall in the bin may be doing so because they are having hadronic interactions with the small amount of material that remains in the bin. Therefore, we check directly neighbouring bins in order to take into account these edge cases and the effective uncertainty on the maps. In an ideal world one would design the bin boundaries such that they are mostly aligned with the material in the GEANT4 model, i.e. our best guess for where material should be, with smaller bins on the order of the expected resolution of the GEANT4 model placed on material edges which can be used to make a more accurate neighbouring bins step. Unfortunately, due to time constraints in the analysis it was not possible to re-derive the maps with custom bin boundaries.

The third step, referred to as the uncertainty veto, is to take into account the DV vertexing position uncertainty. Both the VSI and fuzzy DVs have an associated uncertainty on the vertex position. This is given by the vertexing algorithms in terms of the six non-degenerate elements of the position covariance matrix, which represent the variance in σ_x , σ_y , and σ_z , as well as the covariance of σ_{xy} , σ_{yz} and σ_{xz} . VSI constructs this from the fitting uncertainty whereas fuzzy vertexing constructs this covariance matrix based on the positions of the vertex seeds. One can construct

from this covariance matrix a 3D ellipsoid in the fiducial volume that represents the 1σ uncertainty on the vertex position. In order to take into account this vertexing uncertainty in the veto, to be conservative to within 1σ uncertainty one would then ideally veto the DV as being outside material if any bins of the data-derived map that overlap with the 3D ellipse volume are deemed to be inside material.

It is mathematically and computationally non-trivial to determine the collision between a generic ellipsoid and a wedge-shaped bin. As the 1σ uncertainty is more of a guiding rule-of-thumb rather than a strict boundary, instead the collision between the wedge-shaped bin and the "aligned box" around the ellipse will be tested for collisions with bins. The aligned box is the rectangular prism that is aligned with the ellipsoid's major axes and circumscribes it. One can also define for other purposes a "coordinate bounding box" around the ellipse, which is a rectangular prism aligned with the coordinate axes rather than that of the ellipsoid, which circumscribes the ellipsoid. A diagram of these boxes is shown in Figure 4.6.




Figure 4.6: An uncertainty ellipsoid is shown with its aligned box in red and coordinate bounding box in blue. The major axes of the ellipse shown as black arrows are in the direction of the eigenvectors of the covariance matrix, and of length of the square-root of the corresponding eigenvalue.

If one defines the covariance matrix

$$\underline{\underline{\Sigma}} \equiv \begin{bmatrix} \sigma_x^2 & \sigma_{xy}^2 & \sigma_{xz}^2 \\ \sigma_{xy}^2 & \sigma_y^2 & \sigma_{yz}^2 \\ \sigma_{xz}^2 & \sigma_{yz}^2 & \sigma_z^2 \end{bmatrix}, \quad (4.1)$$

noting that the off-diagonal terms may be negative, then the representation of the ellipsoid in space is

$$\begin{bmatrix} (x - x_0) & (y - y_0) & (z - z_0) \end{bmatrix} \underline{\underline{\Sigma}}^{-1} \begin{bmatrix} (x - x_0) \\ (y - y_0) \\ (z - z_0) \end{bmatrix} = s \quad (4.2)$$

where $s = 1$ for the 1σ contour and $\underline{r}_0 = (x_0, y_0, z_0)$ is the position of the DV. The eigenvectors of $\underline{\underline{\Sigma}}$ form the directions of the major axes of the ellipsoid, with their respective eigenvalues dictating how stretched the ellipsoid is along those axes. Hence, the aligned box is defined by the eigenvalues and eigenvectors, whereas the coordinate bounding box is simply defined by $(\sigma_x, \sigma_y, \sigma_z)$. The normals of the aligned box will also therefore simply be the eigenvectors [82].

The algorithm goes as follows. First, if the DV has already been vetoed by the loose veto, it is not considered in order to save on computation time. Second, if the overall coordinate bounding-box of the ellipse is too large (i.e. if the variance in x , y , or z is too large) the DV is automatically vetoed. This is to save on computational power, as if the ellipsoid is too large it would inevitably intersect with detector material, and it is likely improperly fitted anyways. This threshold was chosen to be $(\sigma_x = 100 \text{ mm}, \sigma_y = 100 \text{ mm}, \sigma_z = 100 \text{ mm})$. Third, each sub-map of the data-derived map is checked, where only bins lying in the overall coordinate bounding box of the ellipsoid that are deemed inside material are considered. The ellipsoid aligned box is then tested for collision detection with every eligible bin.

This collision detection algorithm is the same as that used in video games, known as the separating axis theorem [83]. Public libraries exist with implementations of this algorithm due to its widespread usage. However, in order for it to function efficiently within the analysis framework with as minimal overhead as possible, these libraries were not used and code for the collision detection algorithm was instead written from scratch by the author. If we approximate the wedged-shape bin as a hexahedron with

flat planes connecting the 8 extrema points in (R_{xy}, ϕ, z) , then we are effectively doing collision detection between two hexahedrons. One can define a points matrix and a plane matrix for each of the solids. The points matrix represent the 8 extrema points of the hexahedron and is represented as

$$\underline{\underline{Q}} = \begin{bmatrix} x_1 & x_2 & \cdots & x_8 \\ y_1 & y_2 & \cdots & y_8 \\ z_1 & z_2 & \cdots & z_8 \\ 1 & 1 & \cdots & 1 \end{bmatrix} \quad (4.3)$$

where points (x_1, y_1, z_1) through (x_8, y_8, z_8) are the extrema points. The ones row is for an algorithmic purpose as will be clarified later. As for the planes matrix, these represent the 6 hexahedron planes as

$$\underline{\underline{S}} = \begin{bmatrix} A_1 & B_1 & C_1 & D_1 \\ A_2 & B_2 & C_2 & D_2 \\ \vdots & \vdots & \vdots & \vdots \\ A_6 & B_6 & C_6 & D_6 \end{bmatrix}. \quad (4.4)$$

The planes are represented by equations $A_i x + B_i y + C_i z + D_i = 0$, where $\underline{n}_i = (A_i, B_i, C_i)$ is the normal of the plane pointing outwards from the hexahedron, and $D_i = -\underline{n}_i \cdot \underline{r}_i$, where \underline{r}_i represents a point (x_i, y_i, z_i) that lies on the plane. Note that if one multiplies a row of the $\underline{\underline{S}}$ by a column of the $\underline{\underline{Q}}$ matrix, we get a scaled distance from that particular plane to that particular point, where the distance is positive if the point lies on the "outside the hexahedron" side and vice-versa for inside [83].


Collision detection works by finding if there exists a plane of separation between the two solids. In the simplest case, one can check every plane of both solids, and see if all of the points of the other solid lie on the "outside" of that plane. If that is the case, then that plane is a plane of separation between the two solids and this proves that they do not intersect [83]. This is illustrated in Figure 4.7a and b. Additionally, if there is at least a point of the other solid that has a negative distance with all planes of the solid in question, this means those points lie inside the solid and one can guarantee an intersection.

Therefore, the algorithm first takes the planes of the DV ellipsoid and the points of the bin and multiplies them as $\underline{S}_{DV} \cdot \underline{Q}_{bin}$. This results in a 6 by 8 matrix. If there is at least one row where all 8 entries are positive, this means that there exists a plane of separation between the solids and they do not intersect. If not, the test is inconclusive. If there is at least one column where all 6 entries are negative, this means that at least one point of the bin lies inside the ellipsoid aligned box and we can guarantee an intersection. If not, the test is inconclusive. If the test is still inconclusive, one repeats the above algorithm but using the ellipsoid points and the bin planes.


If the tests still remain inconclusive, there is one last remaining test that needs to be performed. In 3D, one can have a case as illustrated in Figure 4.7c, where the solids are non-intersecting, but no plane-point test using faces of the solids as the planes gives an intersecting result due to their arrangement. In these cases, the plane of separation would be that defined by the cross-product of the edge of one solid and the edge of the other. In the same manner as above, one then checks simultaneously if all the points of one solid lie on one side of the new plane, and all of the points of the other solid lie on the other. If this is the case, a plane of separation has been found and an intersect can be guaranteed. If all of the tests failed, then the solids do not intersect. All edge pairs must be tested [83].

A comparison of the performance of the four separate vetoes in VSI vertexing in Figure 4.8a shows that the data and MC vetoes are on similar footing and yield similar results. It appears around 25% of the DVs that the VSI algorithm reconstructs lie outside our fiducial region. The strict vetoes tend to veto approximately 10% more DVs as the loose veto. The majority of this new vetoing power comes from the vetoing of neighbouring bins, with about 3% coming from the uncertainty-based veto. Effectively no DVs get vetoed by the GEANT4 map veto, and only around 0.2% of DVs have a problematically large covariance.


A comparison four vetoes in fuzzy vertexing in Figure 4.8b also shows that the data and MC vetoes are on similar footing and yield similar results. Here only about 1% of DVs reconstructed by the fuzzy vertexing algorithm lie outside the fiducial volume. This is because the fiducial volume was included as a partial restriction during vertexing. The strict vetoes tend to veto approximately 6% more DVs as the loose veto. Like with VSI, the majority of this new vetoing power comes from the vetoing of neighbouring bins, with about 3% coming from the uncertainty-based veto. Effectively no DVs get vetoed by the GEANT4 map veto or have a problematically large covariance. Significantly more fuzzy DVs pass the veto than in VSI. This is



(a) If a plane can be drawn between the two hexahedrons they do not intersect.




(b) The easiest separating planes to check are the planes of one of the hexahedrons.




(c) There is a case where none of the hexahedron planes are separating, but the hexahedrons do not intersect, that needs to be checked.

Figure 4.7: Diagrams illustrating the separating axis theorem.

because the fuzzy vertexing algorithm tends to have a significantly higher efficiency of reconstructing noise or pileup-related vertices inside the beampipe, which are outside material.



(a) Breakdown for VSI DVs



(b) Breakdown for Fuzzy DVs

Figure 4.8: Comparison of the MC and data, loose and strict vetoes in VSI and Fuzzy DVs.

4.3 Determining the Signal Regions

As with the DV+Jets analysis, there are selections applied to events, namely regarding MET and requirement on numbers of VSI or fuzzy DVs. Both DV types will also have selections, as well as selections on their individual tracks. Similarly to the DV+Jets analysis, selections were optimized in order to maximize signal yield while reducing the number of expected background events in the signal region to about one. This is again due to the large expected uncertainties in the background estimation process.

The event selection for the signal regions has been described in Table 4.1. The following section outlines the development of the track and DV selections for the three VSI and fuzzy DV SRs. Note that these are preliminary selections before background estimation studies were performed, and will likely need to be re-optimized once background estimates are obtained.

4.3.1 Fuzzy DV Signal Regions Selections

Fuzzy DV Signal Regions DV Selections

As there is the 1 Fuzzy DV and 2+ Fuzzy DV SRs, efforts were made to harmonize track selections between the two SRs such that the definition of a "fuzzy DV" is consistent within the analysis. Additionally, efforts were also made to harmonize the DV selections between the two SRs, with the exceptions of the number of tracks and invariant mass, allowing for a more straightforward way of ensuring orthogonality. The fuzzy DV selections are summarised in Table 4.2.

The fuzzy DV pre-selection was developed by Risa Ushioda with a modification on the seeds requirement by the author in order to suit the 2+ Fuzzy DV SR, as well as with the suggestion of the "maximum p_T fraction" cut by the author to reduce background yields for both SRs. The pre-selection keeps the fiducial volume and distance cuts from the DV+Jets analysis, however removes the χ^2 and core tracks cuts as these are no longer relevant for fuzzy DVs.

The "maximum p_T fraction" cut was originally proposed by Oleg Brandt and the author for the DV+Jets analysis but not implemented due to time constraints. Occasionally high p_T tracks can be drastically mismeasured, sometimes to the point of

Fuzzy DV Pre-selection:

- $r_{DV} < 300$ mm and $|z_{DV}| < 300$ mm
- $r_{xy,DV} - r_{xy,PV} > 4$ mm
- $\frac{\max(p_{T,trk})}{\sum_{trks} p_{T,trk}} < 0.9$ after Fuzzy DV track cleanings [Section 4.3.1]
- $N_{seeds} \geq N_{trks}$ if $N_{trks} < 5$. $N_{seeds} \geq 5$ if $N_{trks} \geq 5$ after Fuzzy DV track cleanings [Section 4.3.1]

1 Fuzzy DV SR Full selection:

- DV lies outside detector material as per new material map veto (Section 4.2)
- $N_{trks} \geq 5$, after Fuzzy DV track cleanings [Section 4.3.1]
- $m_{DV} \geq 10$ GeV, after Fuzzy DV track cleanings [Section 4.3.1]

2+ Fuzzy DV SR Full selection:

- DV lies outside detector material as per new material map veto (Section 4.2)
- $N_{trks} \geq 3$, after Fuzzy DV track cleanings [Section 4.3.1]

Table 4.2: DV+MET Fuzzy DV SR Selections

having an obviously unphysically high p_T . Unfortunately, this mismeasurement can result in a DV ending up in the high mass signal region by random chance when it should not. In order to ensure that these DVs do not reach the SR, a cut is imposed to ensure that no track in the DV has an unusually higher p_T than the other tracks. This is done by comparing every track p_T with the scalar sum of all of the track p_T 's, and ensuring that the fraction for any track does not surpass a certain threshold. A distribution of all DVs passing the 2+ Fuzzy DV track cleanings and selections except for this variable are shown in Figure 4.9 (where the 2+ Fuzzy DV SR DV selection is looser and inclusive of the 1 Fuzzy DV SR DV selections), where background can be seen to dominate the higher fraction region. The threshold for the selection was chosen to be 90% as it provided good background rejection while not significantly reducing signal.




Figure 4.9: The fraction of maximum track p_T over sum of track p_T 's plotted for DVs passing the DV selection for the 2+ Fuzzy DV SR. "HP" represents Higgs Portal signal and "WB" represents Wino-Bino signal. The number on the right of the divider represents lifetime whether in τ or $c\tau$ depending on the signal. The first number on the left represents LLP mass. The second number on the left if present represents the mass of the invisible decay product of the LLP.

The DVs are formed via the combination of two-track seeds. Following track cleanings, the number of remaining seeds that still have both tracks in the DV needs

to be recalculated. A higher number of seeds indicates a higher connectivity between the tracks. In theory the minimum number of seeds could be 0 (where this represents the extreme case of all tracks in the DV divided up into completely disconnected pairs, where only one track from each pair actually passes the track cleaning). The maximum number of seeds would be N_{trks} choose 2 (representing every track in the DV being interconnected). Risa Ushioda had shown that requiring five seeds remaining in the DV was sufficient as a selection to remove background for the 1 Fuzzy DV SR. The author for the 2+ Fuzzy DV SR had determined this requirement to be the number of seeds being equivalent to the number of tracks, in order to ensure on average full minimal connectivity between all tracks. As the higher mass DVs do not contribute significantly in the 2+ Fuzzy DV SR, a compromise was reached such that the number of seeds would need to be at least the number of tracks for $N_{trk} < 5$ DVs, and would need to be at least five for $N_{trk} \geq 5$ DVs. This allows for a consistent requirement across SRs while still maintaining optimal expected background and yields for both SRs.

The 1 Fuzzy DV full selection was studied by Risa Ushioda and has ended up identical to that from DV+Jets, with its $N_{trks} \geq 5$, $m_{DV} \geq 10$ GeV, and outside detector material requirements. The 2+ Fuzzy DV full selection was studied by the author and has ended up identical to that of the Higgs Decay with Multiple Implementations ATLAS analysis [78], where there is no explicit mass requirement on the DV, the requirement on the number of tracks is reduced to $N_{trks} \geq 3$, and the outside detector material requirement remains.

Fuzzy DV Signal Regions Track Cleanings

Several preliminary iterations of the track selections were done by both Risa Ushioda and the author. One important development was Risa Ushioda had discovered an extremely efficient type of track cleaning, namely a change to the usual selection on requiring a minimum dot-product angle between the track and the PV-DV vector. This selection intends to help reject low-mass hadronic interaction DVs and tracks that may be effectively from the PV but slightly misreconstructed. However, it was discovered that although the large hadronic interactions background nearly always are DVs that originate from the hard scatter, the hard scatter is not always identified correctly as the PV, at least in the context of events that pass the DV+MET event selections. The selection was therefore expanded to ensure that the minimum dot-product angle

between the track and the PV-DV vector is tested against every pileup vertex in the event as well instead of just the primary vertex. The background rejection efficiency was able to increase significantly with this change, with only a few percent loss in signal yields. This angular variable will be represented by $\Delta\alpha(trk, V - DV, \forall V \in [PV, PU])$.

The efficiency of this cut will of course depend on the radial region. It was found by the author that for 2+ Fuzzy DV SRs, this cut is unfortunately too strict, particularly on low-mass signals very similar to hadronic interactions. However it was found that when investigating tracks that would be cut by the stricter angular cut, the d_0 -significance distribution actually becomes dominated by signal in high d_0 -significance tracks, as can be seen in Figure 4.10. Therefore a new type of cut was developed where a track would need to meet a minimum $\Delta\alpha(trk, V - DV, \forall V \in [PV, PU])$, except if it had a sufficiently high d_0 -sig that could guarantee its quality as a displaced track. The exact values of where to place the selections were done via trial and error, using guidance from plots of the variable in question in the different radial regions.

Most cuts would need to be tuned to the different radial regions, as the background and signal will have different behaviours in different regions, visible in Figure 4.11. Inside the beam-pipe (approximately $R_{xy} < 22$ mm) where there is no material, the background dominates with b -decays. The distributions are simply exponentials, where the background tends to have a lower lifetime than most signals. This radial region for studies was expanded to include up to $R_{xy} < 34$ mm as the region containing the IBL, the closest pixel layer to the beampipe, followed a similar distribution. In the rest of the detector, the distribution shape is dominated by effects of detector material. This region was split at $R_{xy} = 150$ mm such that the fiducial region is divided into three total regions. The upper two regions can have different signals dominating due to different lifetimes.

A maximum allowed $\Delta\phi$ between the track and the PV-DV vector is also imposed with the intention of removing backwards-pointing tracks. However, for heavy LLPs, the boost is low and one should in fact expect occasional backwards-going tracks. For this reason the selection was not set at $\pi/2$, but was allowed to vary as needed in the different radial regions. The exact locations of this cut for the different radial regions was also determined empirically.

An additional cut was done on track p_T much like in the DV+Jets analysis. However the cut had to be reduced from a minimum of 2 GeV to a minimum of 1.1 GeV in order to not be too detrimental to Wino-Bino signal yields. Although due to the recon-

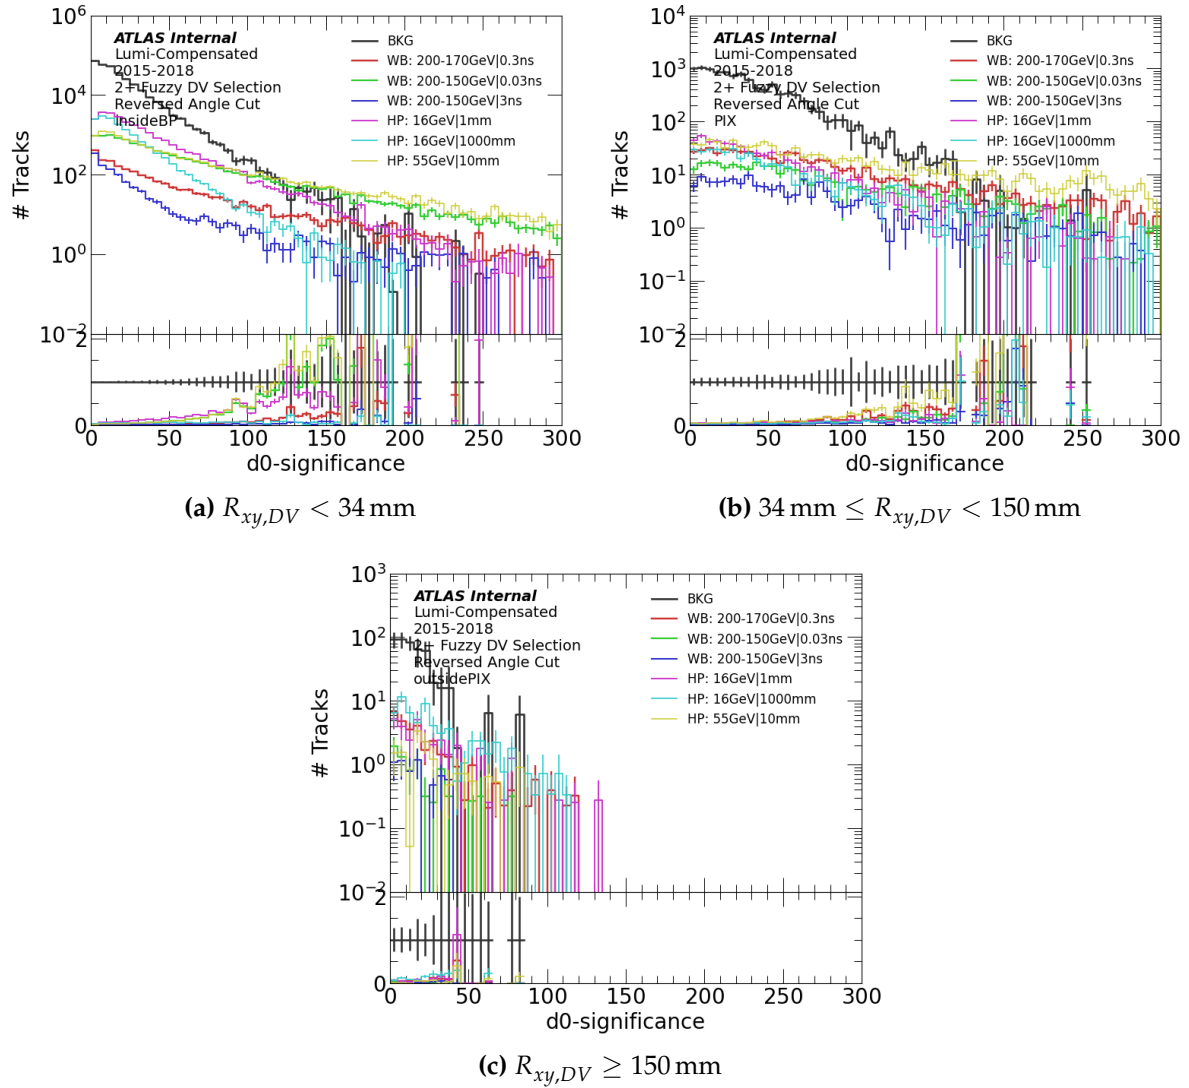


Figure 4.10: Distribution of background and signal track d_0 -significance in the different radial regions. This has the track cleaning applied, except that the minimal angle cut is reverse and the d_0 -significance condition is not applied. This therefore represents tracks that would have been removed by the minimal angle cut if the d_0 -significance condition were not applied.




Figure 4.11: Distribution of background and signal DVs with respect to the DV R_{xy} with the material map veto applied.

struction algorithm, no tracks have $p_T < 1$ GeV, this very minor cut still is important in fine-tuning the background rejection and signal yield optimization.

The final track cleanings, summarised in Table 4.3, were developed by the author after several rounds of iteration with Risa Ushioda to ensure optimal performance for both SRs.

Fuzzy DV Signal Regions Yields

From MC background, the 1 Fuzzy DV SR is expected to see 1.74 ± 2.15 background events, whereas the 2+ Fuzzy DV SR is expected to see 0.39 ± 0.39 background events. The large uncertainties are due to a very low number of raw events in the MC in the signal region.

The Higgs Portal model has sensitivity in both the 1 Fuzzy DV and 2+ Fuzzy DV SRs, with yields shown in Figure 4.12a and b respectively. There is more sensitivity for the mid-range lifetimes from $c\tau = 1$ mm to $c\tau = 100$ mm in the 2+ Fuzzy DV SR,

Track Cleanings

- All hits must $R_{xy,hit} \geq R_{xy,DV}$
 - Track $p_T > 1.1 \text{ GeV}$
-

For DVs with $R_{xy,DV} < 34 \text{ mm}$:

- $\Delta\alpha(trk, V - DV, \forall V \in [PV, PU]) > 0.15$ **except if** track $d_0\text{-sig} > 60$.
 - $\Delta\phi(trk, PV - DV) < 1.75$
-

For DVs with $34 \text{ mm} < R_{xy,DV} \leq 150 \text{ mm}$:

- $\Delta\alpha(trk, V - DV, \forall V \in [PV, PU]) > 0.15$ **except if** track $d_0\text{-sig} > 70$.
 - $\Delta\phi(trk, PV - DV) < 1$
-

For DVs with $R_{xy,DV} \geq 150 \text{ mm}$:

- $\Delta\alpha(trk, V - DV, \forall V \in [PV, PU]) > 0.05$ **except if** track $d_0\text{-sig} > 5$.
 - $\Delta\phi(trk, PV - DV) < 1$
-

Table 4.3: DV+MET Fuzzy DV Track Cleanings

however there is a slightly better sensitivity in the 1 Fuzzy DV SR for the highest lifetime, due to the fact that those cases are less likely to have both DVs decaying in the fiducial volume.

There is effectively no sensitivity for LLP scalars of 5 GeV mass for any lifetime. This is because these decays look virtually identical to b -decays and there are no variables found that can be used to discriminate between the two. There appears to be the strongest sensitivity at $c\tau = 10$ mm, with sensitivity increasing with LLP scalar mass. $c\tau = 1$ mm and $c\tau = 100$ mm points share similar sensitivity, with sensitivity tailing off at $c\tau = 1000$ mm. A dedicated sensitivity study of how low of a Higgs branching ratio can be excluded with this analysis compared to existing limits using these yields is performed in the following Section 4.4. As for discovery potential, assuming a three standard deviation difference from the estimate needed for evidence of discovery, and assuming as a very crude conservative approximation of the background of being 1 ± 1 roughly based on the background MC yields and statistical uncertainty of 0.39 ± 0.39 , one could have evidence of discovery with only observing approximately 4 events. This means that the most sensitive point of $m_S = 55$ GeV and $c\tau = 10$ mm could have evidence of discovery if the branching ratio of the Higgs to the scalar was larger than approximately 0.5%.




Figure 4.12: Yields for the Higgs Portal model in the Fuzzy DV SRs.

As for the Wino-Bino model, there is no previous known search of this particular model. Yields are shown in both the 1 Fuzzy DV and 2+ Fuzzy DV SRs, with yields shown in Figure 4.13a and b respectively for a $\Delta m_{\tilde{\chi}_2^0, \tilde{\chi}_1^0} = 40$ GeV. The model has its greatest sensitivity in the 1 Fuzzy DV SR, however it notably also has some level of sensitivity in the 2+ Fuzzy DV SR despite in theory only one DV being present in the Feynman diagram. This is in most cases likely due to the $b\bar{b}$ decay being reconstructed as two smaller vertices instead of one large fuzzy vertex. Sensitivity is highest for $\tau = 0.03$ ns and for the lowest masses.

As a rule of thumb, a difference of two standard deviations from the estimate (95% confidence) is needed to be observed for exclusion, three standard deviations is needed for evidence of discovery, and five standard deviations is needed to claim discovery. For the 1 Fuzzy DV SR, using the background MC yields and statistical uncertainty directly as a crude estimate, this means that there is exclusion power for points with yields greater than about 6 events, and discovery potential for points with yields greater than about 8 events. This means the analysis will likely be able to exclude up to 750 GeV of neutralino mass at 0.03 ns lifetime, and up to 450 GeV of neutralino mass at the higher 3 ns of lifetime for this case of $\Delta m_{\tilde{\chi}_2^0, \tilde{\chi}_1^0} = 40$ GeV. Similarly the analysis would be likely able to have potential for evidence of discovery for 700 GeV of neutralino mass at 0.03 ns lifetime, and up to 400 GeV of neutralino mass at the higher 3 ns of lifetime for the same case of $\Delta m_{\tilde{\chi}_2^0, \tilde{\chi}_1^0} = 40$ GeV.

4.3.2 VSI DV Signal Region Selections

VSI DV Signal Region DV Selections

The VSI DV selections were studied by Emily Filmer and ended up remaining identical to the DV+Jets DV selections. These are summarised in Table 4.4.

VSI DV Signal Region Track Cleanings

These track cleanings were initially developed by the author and later finalized through studies by Evan Michael Carlson, particularly with an improvement in the selection involving the maximum allowed track angle with respect to the PV-DV vector. As the fiducial volume can be roughly divided into three radial regions, and there are both

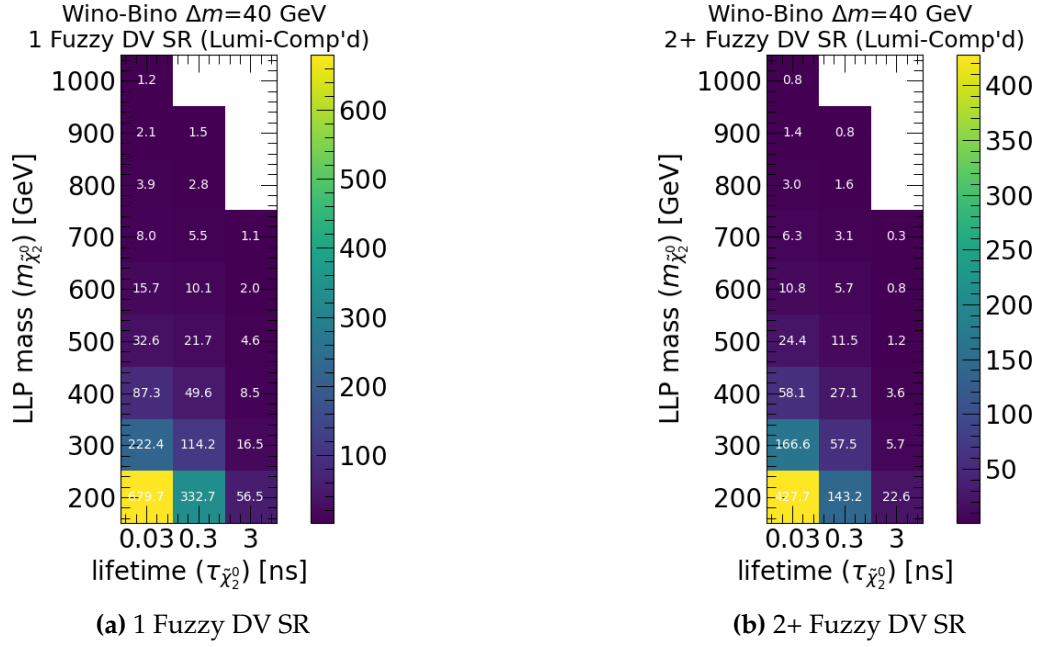


Figure 4.13: Yields for the Wino Bino model in the Fuzzy DV SRs for the case of $\Delta m_{\tilde{\chi}_2^0, \tilde{\chi}_1^0} = 40$ GeV.

Pre-selection:

- $r_{DV} < 300$ mm and $|z_{DV}| < 300$ mm
- $r_{xy,DV} - r_{xy,PV} > 4$ mm
- $\chi^2/\text{DOF} < 5$
- $N_{trks,core} \geq 2$, after track cleanings [Section 3.2.4]

Full selection:

- DV lies outside detector material as per new material map veto (Section 4.2)
- $N_{trks} \geq 5$, after VSI DV track cleanings [Section 4.3.2]
- $m_{DV} \geq 10$ GeV, after VSI DV track cleanings [Section 4.3.2]

Table 4.4: DV+MET VSI DV Selections

core and attached tracks, each variable considered can be plotted in all six regions to determine the optimal location for the selection. Tracks were additionally classified as "good" (descends from the primary OSMLLP of the DV), "bad" (is a crossed track), or "unknown" tracks, in order to ensure that cuts were not made in such a way to select specifically on pathological signal tracks. An example of such plots are shown in Figure 4.14, which shows the distribution of the dot-product angle between the track and the PV-DV vector.

Following preliminary cuts guided by the plots, the cuts were manually moved to empirically optimised for low background and high signal yields. The final track cleanings are summarised in Table 4.5, where $R_{xy,BP} = 22$ mm represents the inner radius of the beampipe.

The selection on the radius of the track's first hit $R_{xy,hit}$ needing to be at a larger radius than the radius of the DV $R_{xy,DV}$, is taken from DV+Jets analysis. The p_T and d_0 -significance cut were loosened to be $p_T \geq 2$ GeV and $d_0\text{-sig} \geq 10$ for all track types and radial regions, which was permissible due to the lower backgrounds in this analysis. It was also found that the cuts on the minimum $\Delta\alpha(trk, PV - DV)$ would only need to be applied inside the beampipe without an addition p_T requirement, and that the maximum $\Delta\alpha(trk, PV - DV)$ requirement on attached tracks would only need to be applied outside the beampipe and could be extended to 2.5 from $\pi/2$. The motivation for this extension is that the heavy LLPs that are being targeted will not have a high boost and do in fact have a chance of having backwards tracks. Both these angular cuts can be justified through the plots in Figure 4.14.

VSI DV Signal Region Yields

From MC background, the selections are strict enough that in fact no events pass the selection. The R-Hadron signal was also used for VSI DVs in the early Run 2 DV+MET analysis, allowing for a direct comparison in sensitivity [84]. To allow for direct comparisons with the paper plots, yields for a fixed $m_{\tilde{\chi}_1^0} = 100$ GeV and a fixed $\Delta m_{\tilde{g}, \tilde{\chi}_1^0} = 100$ GeV are shown in Figure 4.15a and b respectively. Although the MC background estimate is not substantive, if one as a crude approximation assumes the analysis will observe on the order of 1 ± 1 background events in the SR, one would be able to exclude models up to a 95% confidence if they have yields surpassing 3 events in the signal region, or have evidence of discovery is they surpass 4 events.

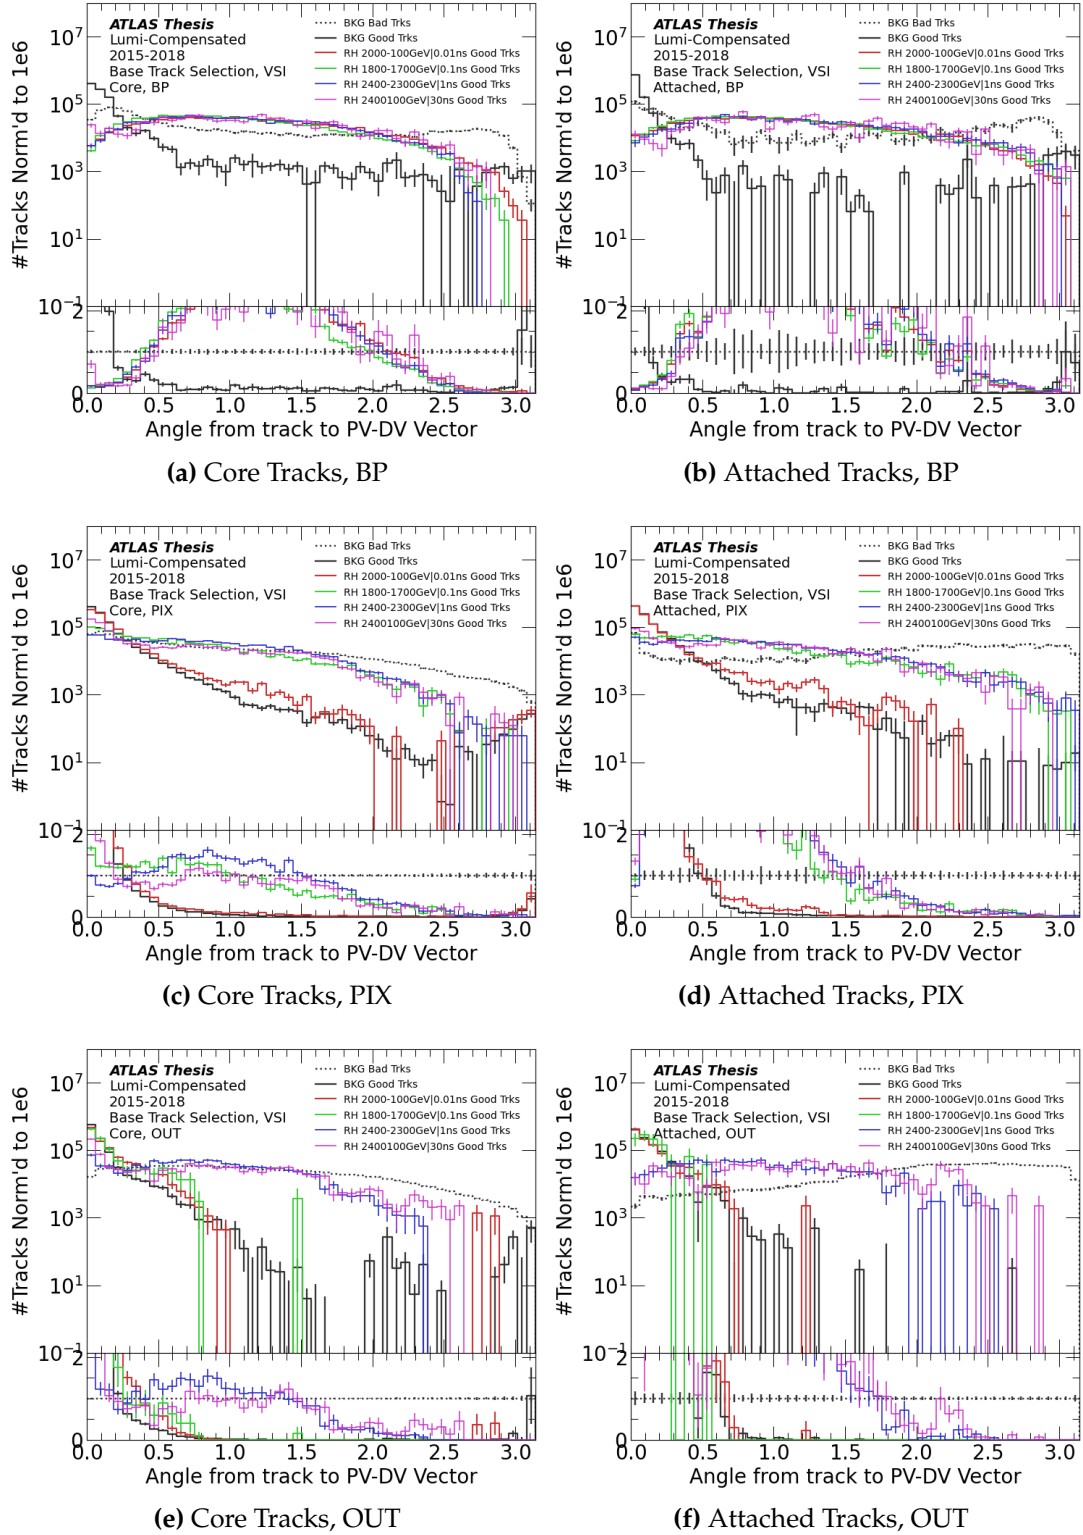


Figure 4.14: Distributions of the angle between tracks and the PV-DV vector for attached and core tracks in the three different radial regions. All tracks from all DVs are included, with the only track selections applied are the two "base" track selections (radius of first hit and backwards track). "BP" represents the $R_{xy,DV} < 22$ mm radial region, "PIX" represents the $22 \text{ mm} \leq R_{xy,DV} < 150$ mm region, and "OUT" represents $R_{xy,DV} \geq 150$ mm.

Track Cleanings

- All hits must $R_{xy,hit} \geq R_{xy,DV}$
 - Track is NOT both $\Delta\phi(trk, PV - DV) > 3$ and $|d_0| < 1$
-
- Track $p_T \geq 2 \text{ GeV}$
 - Track $d_0\text{-sig} \geq 10$
 - $\Delta\alpha(trk, PV - DV) \geq 0.2$ for DVs with $R_{xy,DV} < R_{xy,BP}$
 - $\Delta\alpha(trk, PV - DV) \leq 2.5$ for **attached tracks** in DVs with $R_{xy,DV} \geq R_{xy,BP}$

Table 4.5: DV+MET VSI SR Track Cleanings

The previous analysis excluded in the fixed $m_{\tilde{\chi}_1^0} = 100 \text{ GeV}$ case up to a maximum of $m_{\tilde{g}} = 2350 \text{ GeV}$ at $\tau_{\tilde{g}}$ of about 0.1 to 0.3 ns [84]. In this analysis, from the yields we can roughly expect to exclude up to possibly above $m_{\tilde{g}} = 2550 \text{ GeV}$ or have evidence of discovery up to $m_{\tilde{g}} = 2500 \text{ GeV}$ for the same lifetime range. In the fixed $\Delta m_{\tilde{g}, \tilde{\chi}_1^0} = 100 \text{ GeV}$ case, the previous analysis excluded up to $m_{\tilde{g}} = 1800 \text{ GeV}$ at $\tau_{\tilde{g}}$ of about 0.1 to 0.3 ns [84]. In this analysis we can expect to exclude up to $m_{\tilde{g}} = 2050 \text{ GeV}$ or have evidence of discovery up to $m_{\tilde{g}} = 2000 \text{ GeV}$ in the same lifetime range.

4.4 Higgs Portal Sensitivity Study

This section outlines a sensitivity study for the Higgs Portal model using the yields obtained in the previous section, in order to determine how low in terms of the Higgs branching ratio the analysis is able to exclude. This is done via the CLs formalism [70], using pyHF [71][72]. A workspace was developed to statistically combine both the 1 Fuzzy DV and 2+ Fuzzy DV SRs, taking into account placeholder uncertainties.

For both the signal and background, MC statistical uncertainties and a luminosity error uncertainty of 1.7% (as determined by the collaboration and used in the DV+Jets analysis [1]) are taken into account as such.

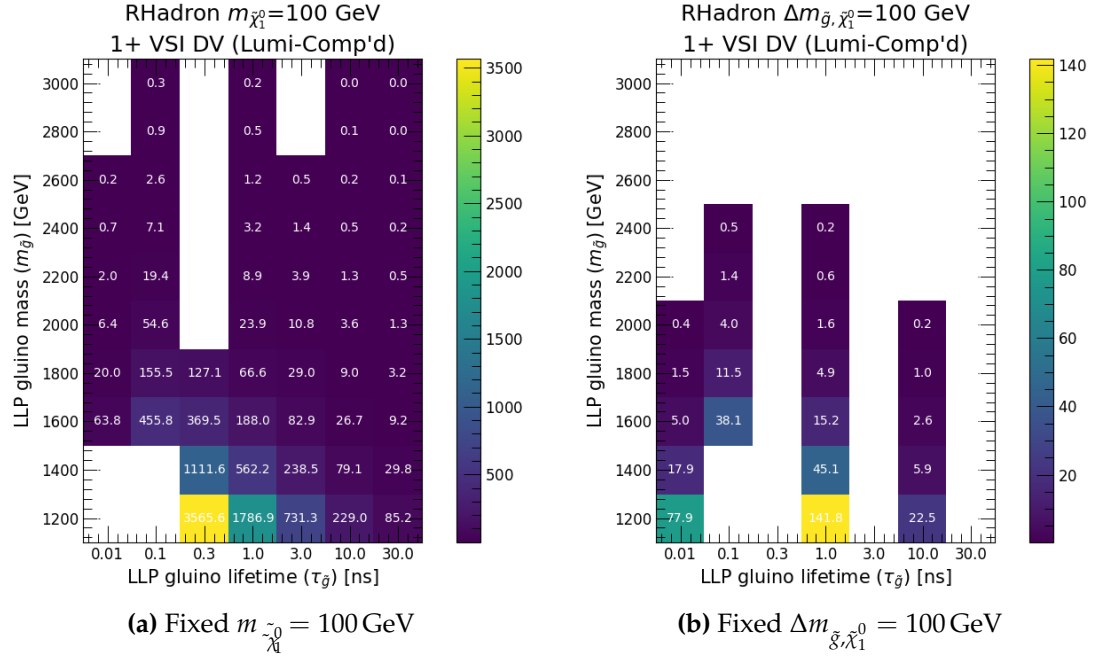


Figure 4.15: Yields for the R-Hadron model in the 1+ VSI DV SR, intended to complement regions plotted in the results in the early Run 2 DV+MET paper to plots [84].

The background has an additional uncorrelated conservative 70% uncertainty to represent the uncertainty that will likely arise from the background estimate process, taking from the uncertainties present in the DV+Jets analysis [1].

The signal has generous correlated uncertainties pertaining to tracking (20%), QCD scale (20%), and PDF and α_s variations (20%), all chosen to be generous with respect to the samples used in DV+Jets that had the highest uncertainties. A MET trigger uncertainty of 1% was also applied, due to the analysis including part the region below the MET trigger plateau, taken from the Higgs to invisible analysis who similarly used this MET turn-on region [85].

The final expected branching ratio exclusion plot is shown in Figure 4.16. Although uncertainties are very large due to the very conservative uncertainties applied in the workspace, this shows that the analysis has the potential to exclude even below 1% branching ratio to a hadronically-decaying LLP scalar of more than 35 GeV mass, which would be the most stringent exclusion to date for this case, compared to the latest ATLAS [74] and CMS [75] results which exclude this down to about 4% at a $c\tau = 10$ mm for the more sensitive higher masses. The $m_S = 55$ GeV and $c\tau = 10$ mm signal point has the highest exclusion power, able to exclude down to 0.05%, which is

an order of magnitude improvement over the current hadronic-decay limits. Note that the current overall limit on the Higgs branching ratio to undetected products is 15% [30].




Figure 4.16: The expected sensitivity for the exclusion of the Higgs branching ratio to LLP scalars that decay hadronically for different LLP masses in different colours. The dotted line represents the expected limit, with the shading representing one σ of uncertainty. Different colours represent different LLP masses. The Higgs to undetected limit [30] is shown with a black line, and the best current ATLAS [74] and CMS [75] limits for hadronically decaying LLP Higgs Portal scalars at $c\tau = 10$ mm is shown with a red star.

4.5 Conclusion and Future Steps

The analysis is currently in progress at the time of this writing, with the main focus of the team being on working on background estimation methods. The background methods are inspired by those used in DV+Jets, however due to issues caused by the different event landscapes from the shift of jets to MET, new issues have arisen that need to be solved, requiring some creative new solutions to make the estimation methods more robust. As the methods being studied are still in flux and not finalized they will not be detailed in this thesis. Following background estimation, signal

region definitions will inevitably need to be re-optimized. However, the promising sensitivity of the analysis can hopefully motivate the importance of seeing this analysis to completion.

Looking forward past this analysis to the future of DV analyses, Run 3 is already underway with a suit of new specialized LLP triggers being developed and investigated. Following Run 3, Run 4 will be with an upgraded LHC designed to deliver higher luminosity. The current inner detector in ATLAS is not suited to deal with such high luminosity. Thus a replacement needs to be created in order to allow for analyses that depend so critically on tracking performance, such as DV analyses, to be able to thrive in this new high-pileup environment. This replacement is the ATLAS Inner Tracker (ITk), and the next chapter will detail the author's contribution to its pre-production and production stages.

Chapter 5

ATLAS Inner Tracker Upgrade (ITk) Strips Sensors Quality Control

“So the ITk is being delayed at the same rate as time...”

— UK ATLAS PhD Student during the 2020 STFC Summer School
at Lancaster, held on March 2021 on Zoom

Following Run 3, a high luminosity upgrade will be installed on the Large Hadron Collider (HL-LHC), giving 4000 fb^{-1} over the course of its planned 10 years of lifetime. The current ATLAS Inner Detector is unable to cope with these conditions, and is approaching the end of its lifetime [86]. A fully-silicon replacement, the ATLAS Inner Tracker (ITk), is therefore currently being developed, comprising both pixel and strip sensors. The strips portion of the ITk comprises 22000 sensors of 8 types: 2 barrel and 6 endcap [87].

Every sensor, with a schematic as shown in Figure 5.1, needs to be evaluated for quality control (QC), which is performed at various institutes, each with their own independently-created test setups with slightly different internal workflows, before they can be integrated into modules that will go into the ATLAS detector. These sensors are similar in detection mechanism to the current ATLAS SCT detecting technology, as outlined in Section 2.2.1.

The decision on whether or not the wafer passes or fails QC should be consistent across the institutes. As the wafers are often transferred between institutes during various stages of production, the test data needs to be kept in a centralized location,




Figure 5.1: a) Bird's-eye view [88] and b) cross-section of a strip sensor

the ITk database (ITk DB) [89]. For these reasons, a common framework with common algorithms is needed to objectively and automatically assign pass/fail decisions to sensors, interface with the common ITk information database, and do reporting. The qualification task is defined as developing software to address this need. The intended workflow of the use of the software developed in this qualification task is described in Section 5.1.

More specifically this task is broken down into several required core aspects that need to be developed. First, functions need to be developed that can read and convert data from the test data files produced by the LabView-based quality control testing setups into a python object. However, this was complicated by the fact that there were multiple competing data formats at the time of development due to it being in the early pre-production stage. Because of this, functions also need to be developed that can convert an old data file into a new format once the new format had been decided on. Functions also need to be developed in order to download test data from the ITk DB into a python object. Functions then need to be made to find tests that correspond to the batch of interest and then organize them to be sent to the analysis

functions. The primary QC tests performed on all sensors are current-voltage (IV), capacitance-voltage (CV), and metrology, whereas current-stability tests on the sensor as a whole and current-capacitance-resistance (ICR) tests on each individual strip are done on a sample. For every test type a custom function needs to be made for extracting parameters from data, such as breakdown voltage or active thickness, and then based on the technical specifications assign a pass/fail decision. Functions then need to be developed to parse and arrange this into a report for the technician, and send this information to plotting functions to make interactive plots in the report. A script interface then needs to be developed such that the technician can override a decision and choose whether or not to upload the test in question. Finally, functions for uploading the test to the ITk DB need to be developed.

As for the context of this task, collaborators working with the ITk DB have made a some generalized functions to help deal with authentication when sending application programming interface (API) requests, which was later superseded by a python library. These were used in order to keep things consistent between other scripts that used the ITk DB. Until this point, technicians have been making individual limited-scope bash scripts to comb through existing data to search for issues. These were not consistent across institutes and only work with the specific data format they were developed for. These were therefore not used and this functionality was redeveloped from scratch. The specifications detailing pass/fail decisions are in the technical specifications document [90]. Some algorithms for extracting parameters needed to be used for pass-fail decisions already existed within the collaboration for the test analysis functions, and where possible, these were adapted to the needs of the task. The rest was developed from scratch. Python 3 was used, and the batch reporting was produced as a local webpage with interactive plots using the graphing library Plotly-JS [91].

The responsibilities on this task were shared with Dominic Jones, with myself taking the primary responsibility for the task. Dominic helped write parts of the analysis functions for the IV test, stability test, and strip test. I was responsible for the rest of the task, namely comprising developing the framework, the interactive reporting interface, the extraction of data from data files, liaising with the ITk DB, the general analysis function code, and specific analysis functions for CV and metrology tests.

5.1 Workflow

Each institute has its own custom setup for performing QC tests. The workflow followed by the institutes is as follows, shown visually in Figure 5.2. A common set of LabVIEW scripts within the collaboration provide an interface to perform the given test with the setup and output a local file in a standard format with the test data.

The goal of the QC scripts is to extract and calculate relevant parameters from these local data files and make an automatic decision on the sensor, in whether or not it passes, fails, is marginal, or if there was likely an error with the measurement.

The ITk DB stores and tracks all information relating to parts of the ITk. The QC scripts are designed to interface with this database such that QC tests, extracted parameters, and decisions can be automatically uploaded. Additionally, the QC scripts are also designed to download raw test data and stored extracted results from the ITk DB in order to do batch reporting. With this, results of a wafer can be compared against its batch with interactive plots that can be used to make decisions on accepting a batch or to investigate issues.

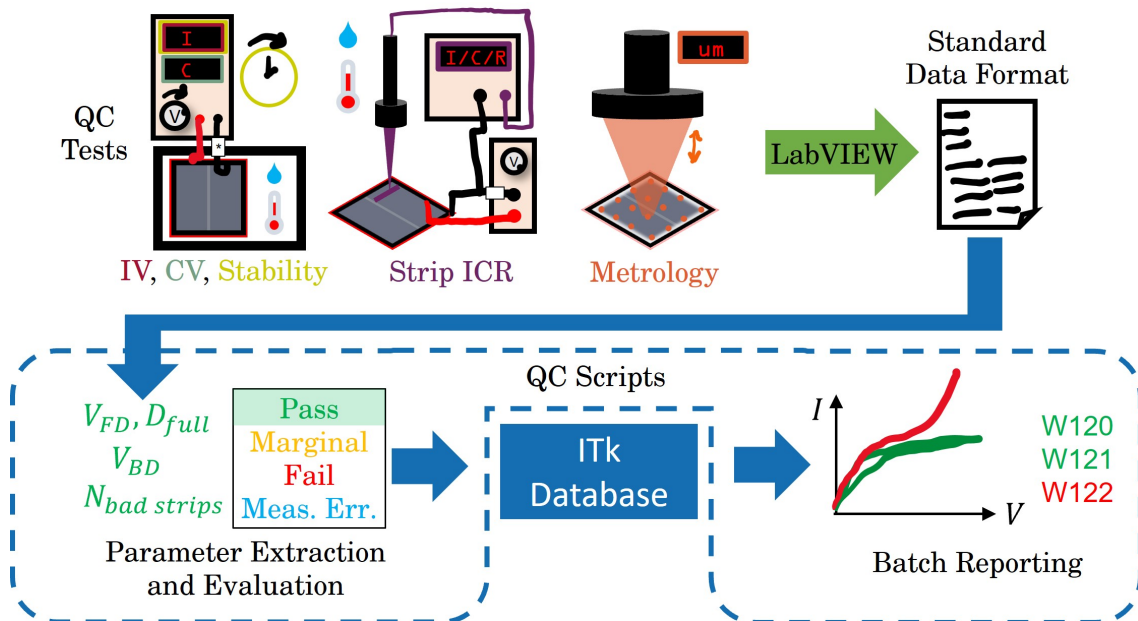


Figure 5.2: Visual representation of the QC workflow.

5.2 An Example of QC Analysis: Treatment of CV Tests

The capacitance-voltage (CV) test is performed by measuring the capacitance of the entire sensor as the high voltage applied across it changes. The technical specifications lists two relevant criteria for sensors to be deemed as passing: 1) the active thickness $D_{full} > 270 \mu\text{m}$ and 2) the initial voltage of full depletion $V_{fd} < 350 \text{ V}$ [92].

As the voltage across the sensor increases, the depletion region slowly grows until it cannot grow further, being limited as it begins to approach the backplane of the sensor. This state at which the depletion region can no longer grow with voltage is known as "full depletion". When a particle passes through the sensor, it is only when it passes through this depletion region where electron-hole pairs are produced, which are picked up by the strip contacts and registered as a hit. Therefore, the active volume of the strip and the sensor is directly related to the thickness of the depletion region. Therefore, the thickness of the depletion region is known as the "active thickness" of the sensor.

This depletion region is where the largest contribution to the capacitance of the sensor originates. As the depletion region is deprived of charge, the positive and negative contacts on either side of the sensor therefore act like a parallel-plate capacitor, as

$$C = \frac{\epsilon_0 \epsilon_r A}{D} \quad (5.1)$$

where C is the capacitance, A is the active area of the sensor, ϵ_0 is the vacuum permittivity, and ϵ_r is the relative permittivity of silicon [93]. As the voltage across the wafer increases the thickness of the depletion region, D increases and C therefore decreases until it hits the plateau region.

The voltage V and capacitance C can be written in terms of x and y , where x is defined as the absolute value of the voltage $x = |V|$, and y is defined as the inverse square of the capacitance $y = 1/C^2$. Parameterizing in this way results in the CV curve acting as a linear increase in y with respect to x until it finally flattens off and stays constant after $x = |V_{fd}|$ [93].

Previous internal rough scripts in the collaboration approached the problem as fitting two individual lines to the data in limited voltage ranges. However, this approach relies heavily on recalibrating the fit ranges for every sensor type, on which

fitting results can heavily depend. Instead, another approach was used, where instead of fitting two separate lines in limited ranges, a functional form comprising both lines and a smooth spline in between is fitted across the entire range. This effectively leaves the "fitting range" to be itself fitted by the curve-fitting algorithm. If one extrapolates the two lines and their intersect lies at $(x, y) = (f, d)$, the functional form is defined as a straight line of variable slope a until $x = f - w$, another straight line of variable slope b (restricted to be close to 0) after $x = f + w$, and a quadratic spline between $x = f - w$ to $x = f + w$. w parameterizes how sharp the transition is. This is illustrated in Figure 5.3




Figure 5.3: A schematic of the CV curve fit. The two linear sections are outlined by green lines which extend to the transition point. The quadratic spline linking the two sections is in orange.

From this, V_{fd} can be naturally extracted from the fit parameters as $f = |V_{fd}|$, as this is the transition point from the expanding depletion region to plateau. The active thickness can be extracted by simply using the parallel-plate capacitor formula to calculate distance, taking the capacitance to be at $y = d$, and therefore $C_{fd} = 1/\sqrt{d}$. This means

$$D_{full} = \epsilon_r \epsilon_0 A \sqrt{d}. \quad (5.2)$$

The scripts obtain the raw capacitance and voltage data points and perform a check to ensure the data is not empty or not all zeroes (which indicate a failed measurement). The capacitance and voltages are converted to the practical variables $x = |V|$ and $y = 1/C^2$, non-finite points are discarded, y is normalised to its mean, and data below

50 V is discarded as it noise-dominated. The full curve function is fitted to data, and then the curve fit parameters are used to calculate V_{fd} and D_{full} . If the data has not already been classified as a measurement error, a pass/fail decision is made based on the technical criteria for these two variables. Histograms and line traces for all sensors in the batch are plotted against each other to help technicians identify batch-anomalies.

5.3 Batch Reporting

QC approval is done on a batch-by-batch basis. This means that the collaboration reserves the right to reject entire batches, even if they may contain some specification-compliant wafers, if there is reason to believe the batch is unreliable. Therefore, even though the scripts make decisions on individual sensors, reports are produced that show interactive summary plots by batch. This allows humans to visually detect outliers such as in Figure 5.4 and potential anomalies in properties across the batch not immediately obvious to an algorithm. The reporting tool is also designed to provide a concise table summary per batch in a single interface to help keep track of what tests have been completed on which sensors as shown in Figure 5.5. Additionally, the scripts provides direct interactive access to both local and database data in Python for studies and investigations.




Figure 5.4: Annotated screenshots of interactive plot from the QC scripts showing depletion voltage histograms with clear specification-compliant outliers.

Wafer	MANUFACTURING18	ATLAS18_IV_TEST_V1	ATLAS18_CV_TEST_V1
W141		ITKDB_MeasErr None 2019-12-02, 15:25 [Temperature Is Empty!] I at 500 V [nA/cm ²]=None Breakdown Voltage[V]=None	ITKDB_Pass None 2019-12-02, 15:51 [] Depletion Voltage [V]=285.83 Active Thickness [um]=304.01 Neff[1e12/cm ³]=4.0168

Figure 5.5: Screenshot of a portion of the report monitoring table. This shows a part of a row of information for wafer 141, where an IV test was performed but was empty due to a measurement error, and a CV test was performed and passed specifications. Having the information in one interface allows technicians to easily spot issues and know how to find the test in question.

5.4 Current Status

The QC scripts have proven to be a robust, reliable and intuitive interface for sensor evaluation, reporting and monitoring. Since their introduction, the library has already processed 3000 sensors through preproduction and production in 7 sites in 9 institutes in 5 countries. As we enter production, the scripts are undergoing continuous development to add new features useful to our QC sites.

Since the completion of the qualification task, the scripts have been taken over by other members of the collaboration who are maintaining them and continuously adding new features to suit its needs.

Chapter 6

Conclusion and Outlook

“At the end of the day, us particle physicists are basically just professional histogrammers.”

— Anja Beck, Professional Histogrammer, LHCb, University of Warwick

Unfortunately, no evidence for new physics was found over the course of this thesis. However, this does not mean that the work done does not have impact. The searches and interpretations presented in this work imply important constraints on what form new physics can take.

Chapter 3 detailed the first search of its kind in Run 2, searching for long-lived particles with displaced vertex signatures in events triggered by multijets. The unique signature allowed for a fairly general model-independent search, however results were interpreted with R-parity violating supersymmetry models. Limits were placed on the existence on 2.4 TeV gluino-produced long-lived neutralinos down to at least 300 GeV for lifetimes between 0.01-1 ns. Limits were also placed on electroweak-produced long-lived neutralinos up to at least 1.2 TeV in the lifetime range of 0.01 to 10 ns. The general nature of the search additionally lends itself well to extending its impact in setting limits to other theories through possible reinterpretations.

Chapter 4 detailed the full Run 2 search of long-lived particles with displaced vertex signatures in events triggered by missing transverse energy. This followed a similar result made with only early Run 2 data from 2016 [2] but was however been completely overhauled with a more robust analysis design and the incorporation of

fuzzy vertexing. Interpreting sensitivity for a R-hadron gluino model where long-lived gluinos decay into quarks and an invisible neutralino, the search has the potential to exclude up to 2050 GeV or have evidence of discovery up to 2000 GeV of gluino mass at 0.1 to 0.3 ns of lifetime, at a gluino-neutralino mass difference of 100 GeV. This improves on the early Run 2 analysis result of excluding 1.8 TeV of gluino mass in the same lifetime range. This analysis will also have the first such result in ATLAS for Wino-Bino models where a long-lived neutralino-2 decays into an invisible neutralino-1 and an off-shell Higgs that decays to a b quark-antiquark pair. Here we have the potential to exclude up to 750 GeV, or have evidence of discovery up to 700 GeV, of neutralino-2 mass for a lifetime of 0.03 ns and mass-splitting of 40 GeV. This analysis has the sensitivity to exclude the branching ratio of the Higgs decaying into hadronically-decaying long-lived particle pairs down to 0.05%, or to have evidence of discovery if the branching ratio is larger than 0.5%, at a lifetime of $c\tau = 10$ mm and long-lived particle mass of 55 GeV. This is an order of magnitude improvement in discovery and exclusion power for hadronically-decaying scalars over the current limits of approximately 5% from ATLAS [74] and CMS [75] for the same point.

Finally, Chapter 5 described the author's qualification task work in the ATLAS inner tracker upgrade (ITk). The development of a software framework for the automatic parameter extraction, evaluation, database interfacing, and reporting of quality control tests on strip sensors was a success, and is currently being used across the collaboration.

There are of course many possible future steps for this line of research. The increase in luminosity expected in Run 3 is already expected to improve results. To more precisely target the Higgs Portal model, sensitivity could be drastically improved by designing a trigger that targets the vector boson fusion (VBF) Higgs production process due to its high cross-section and striking signature. This can be tied together with fuzzy vertexing or jet shape requirements (for example requiring jets have a certain substructure or most of its energy deposited in cells further from the beampipe, which could be possible in the L1Calo) in the trigger, requiring essentially a displaced vertex in the event to allow for more accepting requirements on the VBF jets. Improvements in large-radius tracking in order to improve reconstruction efficiency while maintaining a low rate of reconstructing fake tracks would also benefit sensitivity, as is currently being worked on with the new Release 22 used in ATLAS. There are additionally dozens of other possible long-lived particle signatures in ATLAS that could be targeted that have not yet been fully exploited, such as kinked tracks, or stable charged particles that fully travel outside of ATLAS before decaying. Additionally, as at the moment

most long-lived particle searches target very specific benchmark models. There may be potential benefit in taking inspiration from prompt SUSY and attempting statistical combinations efforts in the field of long-lived particles. Although there may not be a direct increase in discovery potential from such an exercise, the gain comes in the ability to guide and motivate strategic future areas of study. For example, if one notices that the combination tends to poorly exclude a certain sub-class of models with similar signatures, this could motivate developing new triggers and analyses to better target this case.

This is only the beginning. The field of long-lived particle searches is only in its infancy, using only fairly general triggers in a run where it was expected that new physics would have been easier to find. Learning lessons from Run 2, including from this work, the future is bright, with new dedicated triggers and creative strategies to make the most of Run 3 and maximize our discovery potential.

Colophon

This thesis was made in $\text{\LaTeX}2_\epsilon$ using the “hepthesis” class [94].

Bibliography

- [1] 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,
Journal of High Energy Physics **2023** (2023) p. 1,
URL: [https://link.springer.com/article/10.1007/JHEP06\(2023\)200](https://link.springer.com/article/10.1007/JHEP06(2023)200).
- [2] 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** (5 2018) p. 052012,
URL: <https://link.aps.org/doi/10.1103/PhysRevD.97.052012>.
- [3] D. Rouso, D. Jones, P Federicova, B Hommels, A Affolder, *et al.*,
Test and extraction methods for the QC parameters of silicon strip sensors for ATLAS upgrade tracker, **Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment** **1045** (PM2021) (2023) p. 167608, URL: <https://www.sciencedirect.com/science/article/pii/S0168900222009007>.
- [4] 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) p. 1, arXiv: 1207.7214 [hep-ex].
- [5] 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) p. 30, arXiv: 1207.7235 [hep-ex].
- [6] R. Mohapatra, S Antusch, K. Babu, G. Barenboim, M.-C. Chen, *et al.*,
Theory of neutrinos: a white paper, **Reports on Progress in Physics** **70** (2007) p. 1757,
URL: <https://iopscience.iop.org/article/10.1088/0034-4885/70/11/R02>.

- [7] B Gripaios, *Gauge Field Theory*,
University of Cambridge Part III Natural Sciences Tripos (2020),
URL: https://www.hep.phy.cam.ac.uk/~gripaios/gft_lecture_notes.pdf.
- [8] F. Englert and R. Brout, *Broken symmetry and the mass of gauge vector mesons*,
Physical review letters **13** (1964) p. 321.
- [9] P. W. Higgs, *Broken symmetries and the masses of gauge bosons*,
Physical review letters **13** (1964) p. 508.
- [10] P. D. Group, *Review of Particle Physics*, *PTEP* **2022** (2022) p. 083C01,
URL: <https://academic.oup.com/ptep/article/2022/8/083C01/6651666>.
- [11] E. Noether, *Invariant variation problems*, *Transport Theory and Statistical Physics* **1** (1971) p. 186, Translated from Invariante Variationsprobleme, Nachr. d. König. Gesellsch. d. Wiss. zu Göttingen, Math-phys. (1918),
URL: <https://doi.org/10.1080%2F00411457108231446>.
- [12] N. Cabibbo, *Unitary symmetry and leptonic decays*,
Physical Review Letters **10** (1963) p. 531, URL:
<https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.10.531>.
- [13] M. Kobayashi and T. Maskawa,
CP-violation in the renormalizable theory of weak interaction,
Progress of theoretical physics **49** (1973) p. 652,
URL: <https://academic.oup.com/ptp/article/49/2/652/1858101>.
- [14] Z. Maki, M. Nakagawa, and S. Sakata,
Remarks on the unified model of elementary particles,
Progress of Theoretical Physics **28** (1962) p. 870,
URL: <https://academic.oup.com/ptp/article/28/5/870/1858382>.
- [15] T. Sailer, V. Debierre, Z. Harman, F. Heiße, C. König, *et al.*,
Measurement of the bound-electron g-factor difference in coupled ions,
Nature **606** (2022) p. 479.
- [16] 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) 210–255, ISSN: 0146-6410,
URL: <https://www.sciencedirect.com/science/article/pii/S0146641019300109>.

- [17] S. P. MARTIN, *A Supersymmetry Primer*,
Advanced Series on Directions in High Energy Physics (1998) 1–98,
ISSN: 1793-1339, URL: http://dx.doi.org/10.1142/9789812839657_0001.
- [18] G. Hooft,
“Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking”,
Recent developments in gauge theories, Springer, 1980 p. 135.
- [19] V. Faraoni, *Three new roads to the Planck scale*,
American Journal of Physics **85** (2017) p. 865, arXiv: 1705.09749 [gr-qc],
URL: <https://pubs.aip.org/aapt/ajp/article/85/11/865/1057915/Three-new-roads-to-the-Planck-scale>.
- [20] Royal Entomological Society, *Facts and figures*, 2023, URL: <https://www.royensoc.co.uk/understanding-insects/facts-and-figures/>.
- [21] Particle Data Group, *Review of Particle Physics*, *Phys. Rev. D* **98** (3 2018) p. 030001,
URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.030001>.
- [22] E. Corbelli and P. Salucci,
The extended rotation curve and the dark matter halo of M33,
Monthly Notices of the Royal Astronomical Society **311** (2000) p. 441,
ISSN: 0035-8711, eprint: <https://academic.oup.com/mnras/article-pdf/311/2/441/2881340/311-2-441.pdf>,
URL: <https://doi.org/10.1046/j.1365-8711.2000.03075.x>.
- [23] M. Markevitch, A. H. Gonzalez, D. Clowe, A. Vikhlinin, W. Forman, *et al.*,
Direct Constraints on the Dark Matter Self-Interaction Cross Section from the Merging Galaxy Cluster 1E 0657–56, *The Astrophysical Journal* **606** (2004) p. 819,
URL: <https://dx.doi.org/10.1086/383178>.
- [24] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, *et al.*,
Planck 2018 results-VI. Cosmological parameters,
Astronomy & Astrophysics **641** (2020) A6,
URL: https://www.aanda.org/articles/aa/full_html/2020/09/aa33910-18/aa33910-18.html.
- [25] D. G. Cerdeno, *Dark Matter 101: From production to detection*, URL:
https://www.ippp.dur.ac.uk/~dcerdeno/Dark_Matter_Lab_files/DM.pdf.

- [26] ATLAS Collaboration,
Long-lived particle summary plots for Hidden Sector and Dark Photon models,
tech. rep. ATL-PHYS-PUB-2022-007, CERN, 2022,
URL: <https://cds.cern.ch/record/2803320>.
- [27] J. Beacham, *Introduction*, URL: https://indico.cern.ch/event/1216822/contributions/5448388/attachments/2668535/4624988/Beacham_LLP13_Thirteenth_Community_Workshop_Intro_2023June19.pdf.
- [28] P. Coloma, J. Martín-Albo, and S. Urrea, *Discovering Long-lived Particles at DUNE*,
arXiv preprint arXiv:2309.06492 (2023), arXiv: [2309.06492 \[hep-ph\]](https://arxiv.org/abs/2309.06492),
URL: <https://arxiv.org/abs/2309.06492>.
- [29] 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) p. 1, URL:
<https://link.springer.com/article/10.1140/epjc/s10052-021-09411-2>.
- [30] ATLAS Collaboration,
Combined measurements of Higgs boson production and decay using up to 139 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13\text{ TeV}$ collected with the ATLAS experiment,
tech. rep., CERN, 2021, URL: <https://cds.cern.ch/record/2789544>.
- [31] ATLAS Collaboration,
SM Higgs Branching Ratios and Total Decay Widths (update in CERN Report 4 2016),
URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/CERNYellowReportPageBR>.
- [32] 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) p. 1.
- [33] CERN, *SM Higgs production cross sections at $\sqrt{s} = 13\text{-}14\text{ TeV}$ (CERN Report 3)*,
URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/CERNYellowReportPageAt1314TeV2014>.
- [34] C. Schwanenberger, *A review of searches for R-parity-violating SUSY*,
The European Physical Journal C-Particles and Fields **33** (2004) s752,
URL: <https://link.springer.com/article/10.1140/epjcd/s2003-03-813-9>.

- [35] P. Nath, B. Nelson, H. Davoudiasl, B. Dutta, D. Feldman, *et al.*, *The hunt for new physics at the Large Hadron Collider*, **Nuclear Physics B-Proceedings Supplements** **200** (2010) p. 185, URL: <https://www.sciencedirect.com/science/article/pii/S0920563210001490>.
- [36] CERN, *Facts and figures about the LHC*, URL: <https://home.cern/resources/faqs/facts-and-figures-about-lhc>.
- [37] R Garoby, *Upgrade Issues for the CERN Accelerator Complex*, tech. rep. LHC-PROJECT-Report-1110, CERN-LHC-PROJECT-Report-1110, 2008, URL: <https://cds.cern.ch/record/1123676>.
- [38] J. Haffner, *The CERN accelerator complex. Complexe des accélérateurs du CERN*, (2013), General Photo, URL: <http://cds.cern.ch/record/1621894>.
- [39] ATLAS Collaboration, *ATLAS Luminosity Public Results Run 2*, Online (2019), URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>.
- [40] ATLAS Collaboration, *Performance of the ATLAS Trigger System in 2015*, **Eur. Phys. J. C** **77** (2017) p. 317, arXiv: 1611.09661 [hep-ex].
- [41] ATLAS Collaboration, *ATLAS Luminosity Public Results Run 1*, Online (2020), URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults>.
- [42] ATLAS Collaboration, *ATLAS Experiment records “first physics” at new high-energy frontier*, URL: <https://atlas.cern/Updates/Press-Statement/Run3-first-collisions>.
- [43] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, **JINST** **3** (2008) S08003.
- [44] ATLAS Collaboration, *ATLAS insertable B-layer technical design report*, CERN-LHCC-2010-013 (2010), URL: <https://cds.cern.ch/record/1291633>.
- [45] ATLAS Collaboration, *Production and Integration of the ATLAS Insertable B-Layer*, **JINST** **13** (2018) T05008, arXiv: 1803.00844 [physics.ins-det].
- [46] 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*, **PoS EPS-HEP2015** (2015) p. 261, arXiv: 1608.07850 [physics.ins-det].

- [47] W. Walkowiak, *ATLAS Plans for the High-Luminosity LHC*, tech. rep. ATL-PHYS-PROC-2018-048, CERN, 2018, URL: <http://cds.cern.ch/record/2626050>.
- [48] B. Stelzer, A. M. Collaboration, *et al.*, *The new small wheel upgrade project of the ATLAS experiment*, **Nuclear and particle physics proceedings** **273** (2016) p. 1160.
- [49] T. Humanic, *Extracting the hadronization timescale in TeV proton–proton collisions from pion and kaon femtoscopy*, **Journal of Physics G: Nuclear and Particle Physics** **41** (2014) p. 075105, URL: <https://iopscience.iop.org/article/10.1088/0954-3899/41/7/075105>.
- [50] M. Cacciari, G. P. Salam, and G. Soyez, *The anti-kt jet clustering algorithm*, **Journal of High Energy Physics** **2008** (2008) p. 063, URL: <https://iopscience.iop.org/article/10.1088/1126-6708/2008/04/063>.
- [51] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, **The European Physical Journal C** **77** (2017) p. 1, URL: <https://link.springer.com/article/10.1140/epjc/s10052-017-5004-5>.
- [52] 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) p. 1, URL: <https://link.springer.com/article/10.1140/epjc/s10052-018-6288-9>.
- [53] ATLAS Collaboration, *Performance of the reconstruction of large impact parameter tracks in the ATLAS inner detector*, tech. rep., CERN, 2017, URL: <https://cds.cern.ch/record/2275635>.
- [54] J. Long, *The ATLAS Inner Detector tracking trigger at 13 TeV in LHC Run-2 and new developments on standard and unconventional tracking signatures for the upcoming LHC Run-3*, tech. rep., CERN, 2022, URL: <https://cds.cern.ch/record/2813981>.
- [55] ATLAS Collaboration, *Search for long-lived charginos based on a disappearing-track signature in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, **Journal of High Energy Physics** **2018** (2018), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP06\(2018\)022](http://dx.doi.org/10.1007/JHEP06(2018)022).

- [56] ATLAS collaboration,
Search for heavy, long-lived, charged particles with large ionisation energy loss in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS experiment and the full Run 2 dataset, arXiv preprint arXiv:2205.06013 (2022), arXiv: [2205.06013 \[hep-ex\]](https://arxiv.org/abs/2205.06013),
URL: <https://arxiv.org/abs/2205.06013>.
- [57] 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, tech. rep. ATL-COM-PHYS-2019-157, CERN, 2020,
URL: <https://doi.org/10.1103/2Fphysrevd.102.032006>.
- [58] ATLAS Collaboration, *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*, tech. rep., CERN, 2019, URL: <https://cds.cern.ch/record/2669425>.
- [59] J. Goity and M. Sher,
Bounds on $\Delta B = 1$ couplings in the supersymmetric standard model, *Physics Letters B* **346** (1995) p. 69.
- [60] F. Zwirner, *Observable $\Delta B = 2$ transitions without nucleon decay in a minimal supersymmetric extension of the standard model*, *Physics Letters B* **132** (1983) p. 103.
- [61] ATLAS Collaboration,
Search for massive, long-lived particles using multitrack displaced vertices or displaced lepton pairs in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector, *Phys. Rev. D* **92** (7 2015) p. 072004,
URL: <https://link.aps.org/doi/10.1103/PhysRevD.92.072004>.
- [62] D. Rousso, *DV+Jets Brief Updates 2021 11 24*, 2021,
URL: <https://indico.cern.ch/event/1100270/contributions/4628878/attachments/2352385/4013309/Brief%20Updates%202021%2011%2024.pdf>.
- [63] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, *et al.*,
The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *Journal of High Energy Physics* **2014** (2014),
URL: <https://doi.org/10.1007/2Fjhep07%282014%29079>.

- [64] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, *et al.*, *An introduction to PYTHIA 8.2*, *Computer Physics Communications* **191** (2015) p. 159, URL: <https://doi.org/10.1016%2Fj.cpc.2015.01.024>.
- [65] ATLAS Collaboration, *The ATLAS simulation infrastructure*, *The European Physical Journal C* **70** (2010) p. 823, URL: <https://link.springer.com/article/10.1140/epjc/s10052-010-1429-9>.
- [66] GEANT Collaboration, *GEANT4 – A simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) p. 0, URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688>.
- [67] ATLAS Collaboration, *Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector*, tech. rep., CERN, 2015, URL: <https://cds.cern.ch/record/2037702>.
- [68] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, *et al.*, *Particle Decays*, URL: <https://pythia.org/latest-manual/ParticleDecays.html>.
- [69] ATLAS Collaboration, *Early Inner Detector Tracking Performance in the 2015 data at $\sqrt{s} = 13$ TeV*, tech. rep., CERN, 2015, URL: <https://cds.cern.ch/record/2110140>.
- [70] A. L. Read, *Modified frequentist analysis of search results (the CL_s method)*, (2000), URL: <http://cds.cern.ch/record/451614>.
- [71] L. Heinrich, M. Feickert, and G. Stark, *pyhf: v0.7.2*, version 0.7.2, <https://github.com/scikit-hep/pyhf/releases/tag/v0.7.2>, URL: <https://doi.org/10.5281/zenodo.1169739>.
- [72] L. Heinrich, M. Feickert, G. Stark, and K. Cranmer, *pyhf: pure-Python implementation of HistFactory statistical models*, *Journal of Open Source Software* **6** (2021) p. 2823, URL: <https://doi.org/10.21105/joss.02823>.
- [73] CMS Collaboration, *Search for long-lived particles decaying to jets with displaced vertices in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **104** (5 2021) p. 052011, URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.052011>.

- [74] ATLAS Collaboration,
Search for exotic decays of the Higgs boson into long-lived particles in pp collisions at $\sqrt{s} = 13$ TeV using displaced vertices in the ATLAS inner detector,
Journal of High Energy Physics **2021** (2021) p. 1,
URL: [https://link.springer.com/article/10.1007/JHEP11\(2021\)229](https://link.springer.com/article/10.1007/JHEP11(2021)229).
- [75] CMS Collaboration, *Search for long-lived particles produced in association with a Z boson in proton-proton collisions at $\sqrt{s} = 13$ TeV*,
Journal of High Energy Physics **2022** (2022) p. 1,
URL: [https://link.springer.com/article/10.1007/JHEP03\(2022\)160](https://link.springer.com/article/10.1007/JHEP03(2022)160).
- [76] A. Arvanitaki, N. Craig, S. Dimopoulos, and G. Villadoro, *Mini-split*,
Journal of High Energy Physics **2013** (2013) p. 1,
URL: [https://link.springer.com/article/10.1007/JHEP02\(2013\)126](https://link.springer.com/article/10.1007/JHEP02(2013)126).
- [77] K. Rolbiecki and K. Sakurai,
Long-lived bino and wino in supersymmetry with heavy scalars and higgsinos,
Journal of High Energy Physics **2015** (2015) p. 1,
URL: [https://link.springer.com/article/10.1007/JHEP11\(2015\)091](https://link.springer.com/article/10.1007/JHEP11(2015)091).
- [78] B. P. Brau, J. C. Burzynski, M. Danninger, H. Hanif, B. P. Kerridge, *et al.*,
Search for light long-lived particles in pp collisions at $\sqrt{s} = 13$ TeV using displaced vertices in the ATLAS inner detector, tech. rep. ATL-COM-PHYS-2021-885,
CERN, 2021, URL: <https://cds.cern.ch/record/2783923>.
- [79] P. H. Beauchemin, *MetTriggerCalibration*, URL: <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/MetTriggerCalibration>.
- [80] S. Alioli, P. Nason, C. Oleari, and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*,
Journal of High Energy Physics **2010** (2010) p. 1,
URL: [https://link.springer.com/article/10.1007/JHEP06\(2010\)043](https://link.springer.com/article/10.1007/JHEP06(2010)043).
- [81] E. Bothmann, G. Singh Chahal, S. Höche, J. Krause, F. Krauss, *et al.*,
Event generation with Sherpa 2.2, *SciPost Physics* **7** (2019) p. 034,
URL: <https://scipost.org/10.21468/SciPostPhys.7.3.034>.
- [82] A. Wong, SYDE 372: *Introduction to pattern recognition*,
URL: <https://www.eng.uwaterloo.ca/~a28wong/syde372.html>.
- [83] G. Engine, *Advanced vector math*, URL: https://docs.godotengine.org/en/stable/tutorials/math/vectors_advanced.html.

- [84] 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) p. 052012, arXiv: 1710.04901 [hep-ex].
- [85] ATLAS Collaboration,
Combination of searches for invisible decays of the Higgs boson using 139 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment,
Physics Letters B **842** (2023) p. 137963, URL: <https://www.sciencedirect.com/science/article/pii/S0370269323002976>.
- [86] Y. U. et al., *ATLAS17LS – A large-format prototype silicon strip sensor for long-strip barrel section of ATLAS ITk strip detector*, **Nuclear Instruments and Methods in Physics Research Section A** **989** (2021) p. 164928, ISSN: 0168-9002.
- [87] ATLAS Collaboration,
Technical Design Report for the ATLAS Inner Tracker Strip Detector,
tech. rep. CERN-LHCC-2017-005, ATLAS-TDR-025, CERN, 2017,
URL: <https://cds.cern.ch/record/2257755>.
- [88] L.B.A. Hommels et al., *Detailed studies of full-size ATLAS12 sensors*, **Nuclear Instruments and Methods in Physics Research Section A** **831** (2016) p. 167.
- [89] CERN, *ATLAS ITk Production Database*,
URL: <https://itkpd-test.unicorncollege.cz/>.
- [90] The ATLAS/CMS Inner Tracker Upgrade Project,
Technical Specification for the Supply of Silicon Sensors, tech. rep. AT2-IS-CD-0019,
2019, URL: <https://edms.cern.ch/document/2069229/1>.
- [91] Plotly, *Plotly javascript graphing library in JavaScript*,
URL: <https://plotly.com/javascript/>.
- [92] ATLAS ITk Strips Sensors Collaboration, 2016,
URL: <https://indico.cern.ch/event/592256/contributions/2390452/attachments/1388483/2121025/ATLAS17LSTechnicalSpecsR1.3.pdf>.
- [93] C. T. Klein, *Investigation of performance and the influence of environmental conditions on strip detectors for the ATLAS Inner Tracker Upgrade*, (2020),
URL: <https://www.repository.cam.ac.uk/handle/1810/298758>.
- [94] A. Buckley, *The hepthesis L^AT_EX class*,
URL: <https://ctan.org/pkg/hepthesis?lang=en>.

List of figures

1.1	A summary of the fundamental particles of the SM. The fermions have anti-fermion counterparts. Particle masses from [10].	2
1.2	The system of infinite single fermion loops occurring sequentially can be interpreted as a correction to the particle mass.	6
1.3	Feynman diagrams for the various Higgs production process. The two partons on the left-side of the diagram originate from the two different protons being collided.	13
2.1	A diagram of the CERN accelerator complex. The protons start in a hydrogen bottle before being injected into LINAC 2, before finally ending up in the LHC [38].	18
2.2	A diagram of the ATLAS detector [43].	21
2.3	A cross-section diagram of the inner detector of ATLAS [46].	24
3.1	Benchmark models for the DV+Jets analysis. The blue circles represent the DVs, as well as where the RPV coupling λ'' is. The object surrounded by the green dots represent objects that the analysis triggers on, in this case multijets.	37
3.2	Schematic showing the different regions of the analysis in terms of m_{DV} , n_{Tracks}^{DV} , and inside/outside of material, for the MC with no event selection, as well as for the high p_T and trackless data regions. VR represents the validation region, SR represents the signal region.	46
3.3	Visualizations of locations of DVs in $(R_{xy}, \phi, z = 0 \text{ mm})$ that would be considered by the different vetoes as inside material are denoted in red.	49

- 3.4 Legend of colours used to designate types of tracks in schematics in this following chapter. Solid lines represent tracks reconstructed in the DV, whereas dotted lines represent tracks not in the DV. A purple dot represents the PV, and a solid-line open circle would represents the DV. 56
- 3.5 Decay tree of an event showing relevant truth particles. Edges represent particles, vertices are coloured by time distance from the PV (grey if they are at the PV itself). Particles highlighted in bold red are reconstructed in the DV. Dotted lines represent Gen tracks, solid lines represent G4 tracks. PU and fake tracks cannot be truth-matched to the hard scatter and are therefore not included. The particles are highlighted with the same colour share the same OSMLLP. Originating protons of the collision are out of the scope of the plot, but would be above. Note that not all particles involved in the event are included. 58
- 3.6 Truth particle event display of the event shown in Figure 3.5. All truth particles are shown, as well as the beampipe. The particles that have ended up reconstructed into the DV are outlined in red, where it is clear they all originate from the same interaction. Particles sharing the same OSMLLP as the reconstructed tracks are coloured in green. This appears to be a hadronic interaction with the beampipe. 59
- 3.7 Schematic of how the OSMLLP is determined. All tracks highlighted with the same colour share the same OSMLLP. The OSMLLPs are denoted in the schematic. 59
- 3.8 2D histogram of the production and decay times of all truth particles in a sample of events containing DVs with 3 or more tracks in a single DAOD file. The red dotted lines represent the underflow and overflow bins. The cyan lines represent the ct of 1 mm from the collision, where this appears to reasonably split the multimodal distribution into sensible regions. 61
- 3.9 Breakdown of the mass distributions of the various truth-categorizations of DVs, here plotted for $n_{\text{Tracks}}^{\text{DV}}=4$ and inside material. 63
- 3.10 Schematic of the tSP sub-categories. 64

3.11 Breakdown of the mass distributions of the tSP DVs, plotted for $n_{\text{Tracks}}^{\text{DV}} = 5$ inside and outside of detector material. For the inside plot, the two dotted lines are shown in order to guide the eye to demonstrate the inflection in slope in the hadronic interactions curve, here occurring at about $m_{\text{DV}} = 10 \text{ GeV}$	64
3.12 Example of an a) "regular HI" DV and a b) "elastic collinear HI" DV. These plots have the PV-DV vector aligned to the y-axis, whereas the absolute value of the momentum in the direction transverse to the PV-DV vector is plotted in the x-axis. These figures keep an equal aspect ratio in order to accurately reflect the scale of the relative orientations of the tracks. c) shows a schematic of these DVs.	66
3.13 Schematic of the tAX sub-categories.	67
3.14 Breakdown of the mass distributions of the tAX DVs, plotted for $n_{\text{Tracks}}^{\text{DV}} = 4$ inside detector material.	67
3.15 Schematic of the tOther truth categories.	70
3.16 Breakdown of the mass distributions of the tOther DVs, plotted for $n_{\text{Tracks}}^{\text{DV}} = 5+$ inside detector material.	70
3.17 $(\Delta R)_{\text{max}}$ distribution of tSP and non tSP DVs in MC. A cut is selected at < 0.8 to select for a high purity sample of tSP DVs.	74
3.18 The mass distribution of tSP MC is plotted alongside data with a $(\Delta R)_{\text{max}}$ cut, in order to get a proxy for what the tSP components of data with selections look like. The first column is for inside material and the second column is for outside material, with four (first), five (middle) or six or more tracks (last).	75
3.19 tSP MC DVs that only have GEANT4 tracks (in order to isolate the "true HI" DVs that dominate in the high mass from possible decays that only exist in low mass) compared inside and outside of material in order to determine whether or not their shapes are comparable. The DVs have four (left), five (middle) or six or more tracks (right).	77
3.20 Schematic showing the equations of the mass distribution in terms of the primary and secondary exponentials in the different regions.	77

3.21	The estimate in the full MC against the tSP MC as a sort of closure test. The top row is for inside material and the bottom row is for outside material, with four (left), five (middle) or six or more tracks (right). Note that these do not have event weights applied and bin contents are not normalised to bin width. "Nominal" refers to the fit without the secondary component correction.	79
3.22	Visual representation of the procedure for estimating accidental crossings.	80
3.23	Visual representation of a DV and track in the detector, showing the PV-DV vector and the angles $\Delta\alpha$ and $\Delta\phi$	81
3.24	Visual representation of the track momentum constraints. The blue cone represents the constraint from p and p_T . The green cone represents the constraints from p and $\delta\alpha$. The red line represents the PV-DV vector. The two intersections of the cones, in yellow, are the solutions for momenta satisfying all constraints.	82
3.25	Mass distributions in MC divided by subcategory for AX, with what is aimed to be estimated and the two angular estimate methods. These plots are made without event weights.	83
3.26	Comparison of the S^2 distribution from 2 and 3-track vertices that were merged together from the same vs. different events for the trackless SR. The merging portion of the secondary-vertexing algorithm merges vertices with $S^2 < 100$. The deficit in the $S^2 < 100$ distribution for the same-event merging represents vertices that were already incorrectly merged by the algorithm. This plot was produced by Atsushi Mizukami and Emily Anne Thompson, as it appears in the analysis paper [1] . . .	84
3.27	The combined estimate performed in MC is compared to true MC. The difference in the high mass bins is taken to be a non-closure uncertainty for the combined eAX and eMV. The overall systematic for the eAX trace was developed by Emily Anne Thompson. The trace for the eMV estimate is produced by Emily Anne Thompson and Atsushi Mizukami	86
3.28	The blinded validation region in data ($n_{\text{Tracks}}^{\text{DV}} = 4$, $m_{\text{DV}} > 20 \text{ GeV}$), against the combined background estimate, created by the team and taken from the official paper [1].	87

3.29	Event display of the single observed event in the signal regions in data. Created by Moe Wakida. The blue circles represent different proton-proton interactions along the beamline. The green tracks are the reconstructed DV. The shaded brown cylinder represents the beampipe.	91
3.30	The dashed line and the yellow band represents the expected limits and its $\pm 1\sigma$ uncertainty. The solid red line represents the observed limit, with the dotted red lines representing the theoretical $\pm 1\sigma$ uncertainty. These plots are taken directly from the paper [1], created by the team using respective inputs, using code created by both Stefan Richter and myself.	92
3.31	Expected yields for the Higgs Portal model at 100% Higgs branching ratio inclusive of all Higgs production modes for a limited signal scan of LLP scalar mass and lifetime. Samples used to create this plot were produced with Anna Mullin.	93
4.1	An LLP originating from the PV decays into a $b\bar{b}$ pair, which each decay later into several tracks. Black dotted tracks are not reconstructed by the detector. Orange dotted tracks are not reconstructed into the DV, but orange solid tracks are. The blue circle represents the DV, which is point-like in VSI vertexing and a volume in fuzzy vertexing.	96
4.2	Benchmark models for the DV+MET analysis. The blue circles represent the DVs. The object surrounded by the green dots represent objects that the analysis triggers on, in this case MET.	98
4.3	Visualizations of the data-derived and GEANT4 maps that are used for vetoes. These are shown in $(R_{xy}, \phi, z = 0 \text{ mm})$, with "inside material" denoted in red.	103
4.4	A visualization of the GEANT4 map in the inner radial region with $(R_{xy}, \phi, z = 0 \text{ mm})$. Yellow denotes "inside material". DVs with 5+ tracks from a dijets MC background sample are shown as white crosses overlaid on the map only if they happen to lie in bins the map deems as outside material.	104
4.5	Visualizations of regions in red that the GEANT4 map deems as "inside material" but the data-derived map deems as "outside material".	104

4.6	An uncertainty ellipsoid is shown with its aligned box in red and coordinate bounding box in blue. The major axes of the ellipse shown as black arrows are in the direction of the eigenvectors of the covariance matrix, and of length of the square-root of the corresponding eigenvalue.	106
4.7	Diagrams illustrating the separating axis theorem.	110
4.8	Comparison of the MC and data, loose and strict vetoes in VSI and Fuzzy DVs.	111
4.9	The fraction of maximum track p_T over sum of track p_T 's plotted for DVs passing the DV selection for the 2+ Fuzzy DV SR. "HP" represents Higgs Portal signal and "WB" represents Wino-Bino signal. The number on the right of the divider represents lifetime whether in τ or $c\tau$ depending on the signal. The first number on the left represents LLP mass. The second number on the left if present represents the mass of the invisible decay product of the LLP.	114
4.10	Distribution of background and signal track d_0 -significance in the different radial regions. This has the track cleaning applied, except that the minimal angle cut is reverse and the d_0 -significance condition is not applied. This therefore represents tracks that would have been removed by the minimal angle cut if the d_0 -significance condition were not applied.	117
4.11	Distribution of background and signal DVs with respect to the DV R_{xy} with the material map veto applied.	118
4.12	Yields for the Higgs Portal model in the Fuzzy DV SRs.	120
4.13	Yields for the Wino Bino model in the Fuzzy DV SRs for the case of $\Delta m_{\tilde{\chi}_2^0, \tilde{\chi}_1^0} = 40 \text{ GeV}$	122
4.14	Distributions of the angle between tracks and the PV-DV vector for attached and core tracks in the three different radial regions. All tracks from all DVs are included, with the only track selections applied are the two "base" track selections (radius of first hit and backwards track). "BP" represents the $R_{xy, DV} < 22 \text{ mm}$ radial region, "PIX" represents the $22 \text{ mm} \leq R_{xy, DV} < 150 \text{ mm}$ region, and "OUT" represents $R_{xy, DV} \geq 150 \text{ mm}$	124

4.15	Yields for the R-Hadron model in the 1+ VSI DV SR, intended to complement regions plotted in the results in the early Run 2 DV+MET paper to plots [84].	126
4.16	The expected sensitivity for the exclusion of the Higgs branching ratio to LLP scalars that decay hadronically for different LLP masses in different colours. The dotted line represents the expected limit, with the shading representing one σ of uncertainty. Different colours represent different LLP masses. The Higgs to undetected limit [30] is shown with a black line, and the best current ATLAS [74] and CMS [75] limits for hadronically decaying LLP Higgs Portal scalars at $c\tau = 10$ mm is shown with a red star.	127
5.1	a) Bird's-eye view [88] and b) cross-section of a strip sensor	130
5.2	Visual representation of the QC workflow.	132
5.3	A schematic of the CV curve fit. The two linear sections are outlined by green lines which extend to the transition point. The quadratic spline linking the two sections is in orange.	134
5.4	Annotated screenshots of interactive plot from the QC scripts showing depletion voltage histograms with clear specification-compliant outliers.	135
5.5	Screenshot of a portion of the report monitoring table. This shows a part of a row of information for wafer 141, where an IV test was performed but was empty due to a measurement error, and a CV test was performed and passed specifications. Having the information in one interface allows technicians to easily spot issues and know how to find the test in question.	136

List of tables

3.1	DV+Jets Event Selections	41
3.2	RPV/LL Filter Jet Requirements for the High p_T SR.	42
3.3	RPV/LL Filter Jet Requirements for the Trackless SR.	42
3.4	Final Offline Jet Requirements for the High p_T SR.	43
3.5	Final Offline Jet Requirements for the Trackless SR.	43
3.6	DV+Jets DV Selections	46
3.7	DV+Jets Track Cleanings	50
3.8	This table summarises the definitions chosen to categorize different types of reco-tracks. Using these definitions, there are no reco-tracks that fall outside of these 4 categories.	55
3.9	Summary of the yields in the signal regions as estimated by the combined and inclusive estimates, compared to the final unblinded yields in data. These are the final numbers from the paper, calculated by various members of the analysis team [1].	90
4.1	DV+MET Event Selections	100
4.2	DV+MET Fuzzy DV SR Selections	113
4.3	DV+MET Fuzzy DV Track Cleanings	119
4.4	DV+MET VSI DV Selections	122
4.5	DV+MET VSI SR Track Cleanings	125