A Search for Long-Lived Heavy Neutral Leptons in Events with Displaced Vertices Using the ATLAS Experiment

A Dissertation

Presented to

The Faculty of the Graduate School of Arts and Sciences
Brandeis University

Department of Physics

Gabriella Sciolla, Advisor

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

by

Sagar Vidya Addepalli

May 2024

This dissertation, directed and approved by Sagar Vidya Addepalli's Committee, has been accepted and approved by the Faculty of Brandeis University in partial fulfillment of the requirements for the degree of:

DOCTOR OF PHILOSOPHY

Charles Golden, Interim Dean
Graduate School of Arts and Sciences

Dissertation Committee:
Prof. Gabriella Sciolla, Department of Physics,
Prof. Aram Apyan, Department of Physics,
Dr. Maximilian Emanuel Goblirsch-Kolb, European Organization for Nuclear Research

Acknowledgements

This Ph.D. thesis wouldn't have been possible without the involvement and support from so many people in my life. I decided to become a physicist a long time ago, not knowing what it really meant to be one. Since then, I have been impacted with love and guidance at all stages, and I am so thankful for all of it.

First and foremost, I would like to thank my advisor **Gabriella Sciolla**. I still remember our call before I joined Brandeis when you told me about your vision for how my Ph.D. would look like, and I have achieved everything you promised, and more. I am incredibly grateful for every opportunity, advice, and feedback I have received from you in the last five years. Thank you for not giving up on me when my vision would not align with yours, and always making sure I did what is best for me (even if it took a bit (a lot) of convincing). What I have learnt from you, especially in the last year, made me a more confident and well-rounded scientist, and that is something I will forever keep with me. There is only one person who wanted me to excel in this journey even more than I did, and that is you.

I would like to thank the other members of my thesis defense committee, **Aram Apyan** and **Max Goblirsch** for the thorough review of my thesis and the challenging round of questioning during my defense. Aram, thank you for being my second advisor acting as a sounding board when I needed it. I will cherish our Luigia dinners and I look forward to more (I will order whatever you will)! Max, you have read my thesis more times than I have, making sure I add every missing piece that would make it more complete. I have learnt so much from you in the last two years, ranging from running Athena to understanding the intricacies of limit setting, and acknowledging all of it will take a lot more than a few paragraphs. You will always be an inspiration to me, and I can only hope to become as good a physicist and human being as you one day. I promise to learn how to get better at making money the quick way, and to never check my memory leaks.

I am grateful for my amazing postdocs, **Guglielmo Frattari** and **Dominique Trischuk**.

Tacchino, you have taught me to approach science with practicality rather than idealism, that sometimes good is better than perfect, and that the one who gets work done is the MVP. I will dearly miss the banter that we shared and I hope that you never stop sending me funny reels. Dominique, you have taught me to take ownership of and pride in my work, to lead from the back when the opportunity presents, and that I don't need to be miserable to be successful. I will keep making you buy me coffees and cookies no matter where we are.

In my time at BNL, I had the pleasure of working with some incredible scientists on hardware and analysis. I want to thank **Dr. Dave Lynn** for being an amazing mentor for my ITk work. You taught me how to solve problems at the most fundamental level and approach tasks with a scientific temperament. I also want to thank **Stefania Stucci** for her gentle yet impactful teaching, and also for sharing funny jokes and delicious food with me. Thanks to **Dr. Alessandro Tricoli** for welcoming me into the $H \rightarrow WW^*$ team and giving me the opportunity to work on critical tasks related to such a complex measurement. Thanks to **Laura Bergsten** and **Gaetano Barone** for supervising me in the analysis and setting me up for independence. Gaetano, I appreciate your honest advice at every stage of my Ph.D., and I wish that someday I have at least half as many passports as you have. Thanks to **Jiayi Chen**, **Guy Rosin**, and **Chara Kitsaki** for being such great peers and for sharing tasks when I was occupied with coursework or teaching.

This work wouldn't have been possible without my colleagues in the displaced HNL team, and I would like to thank them for their contributions to all the aspects of the analysis. Thanks to the members of the SUSY displaced leptons team for pushing the LRT lepton development in parallel, accelerating both of our timelines together. I have had the pleasure of interacting with some amazing physicists during my Ph.D. outside the context of my analysis. Thanks to **Johannes Junggeburth** for teaching me everything about muons and muon reconstruction. Thanks to **Haider Abidi** for always giving me sound career advice. Thanks to **Robin Hayes** for always being a skype message away from all kinds of help. Thanks to **Audrey Kvam** for teaching me about muon triggers, and always willing to share gossip over coffee.

The students of the Brandeis hep-ex group have been my strongest support throughout this journey. Thanks to **Laura Bergsten** for your guidance in $H \rightarrow WW^*$. Thanks to **Francesca Capocasa** for teaching me everything about LabVIEW, for the amazing times in Ridge, for teaching me how to make pizza, and for the exciting trips to Hicksville. Thanks to **Jiayi Chen** for always being there to listen to my complaints, for all the car rides whenever I needed them, and for being a great friend always. Thanks to **Prajita Bhattarai** for being the best roommate one could ask, for teaching me to go easier on myself, and for nudging me to be less grumpy and more grateful. Thanks to **Zach Schillaci** for helping ease my move to Geneva. Thanks to **Todd Zenger** for your joyful companionship at BNL and at CERN. Thanks to **Emily Duden** for listening to my late night rants in Waltham, for being someone I can always joke with, and for happily lending your couch. Thanks to **Siddharth Singh** for being such an amazing friend whom I can always trust and rely on.

I also want to thank the Brandeis GSAS administration for making the process of moving from India to BNL to Brandeis to CERN very convenient, especially **Joan Thorne** and **Anne Lazerson**. Special thanks to the **UChicago Tier-3 cluster** for giving my analysis jobs a higher priority in the days leading to my thesis submission, without which I wouldn't have managed to meet the deadlines.

I have made many friends along the way, and my Ph.D. would not have been as pleasant without any of them. Thanks to my wingmates from IITB, **Deepesh Singh**, **Aditya Singh**, **Abhijeet Melkani**, **Sanyam Chouksey**, and **Amey Sharma** for being my support system during my time there. Thanks to **Punit Sharma** for truly being there for me from the beginning to the end. We started this together, and we will go a long way together. Thanks to **Arkya Chatterjee** for always being ready to explore Boston or share your room with me. Thanks to **Andrew Smith** for being my first friend in Geneva and for introducing me to so many amazing people whom I call friends today. Thanks to the *l u n* chat for your constant presence and support, and for never making me feel left out no matter how occupied

I was. Thanks to **Dev**, **Elena**, **Tom**, **Eleanor**, **Antti**, **Kevin**, **Laura**, **Lauren**, **Bruno**, **Ki Ryeong**, **Sylvia**, and **Michelle**. I could never have enough of the CERN lunches, BDP sauna Mondays, trivia Tuesdays, ski trips, movie nights, or any moment that I can share with you all.

To my family, ***amma***, ***nanna***, **Vinit**, and **Kalyani** for all the love. Special shout-out to my *pinnis* who have always cheered for me.

Finally, to **Shivani**, for the unconditional support and understanding throughout this whole thing. I couldn't have asked for more.

ABSTRACT

A Search for Long-Lived Heavy Neutral Leptons in Events with Displaced Vertices Using the ATLAS Experiment

A dissertation presented to the Faculty of the
Graduate School of Arts and Sciences of Brandeis University
Waltham, Massachusetts

By Sagar Vidya Addepalli

The existence of right-handed neutrinos with Majorana masses below the electroweak scale could give rise to the smallness of neutrino masses and matter-antimatter asymmetry in the universe. This analysis searches for long-lived right-handed neutrinos, called heavy neutral leptons (HNLs), produced through a W boson decaying into a muon (electron) and a muon (electron) (anti-)neutrino that will mix with an HNL. The ATLAS reconstruction chain went through significant improvements in preparation for data taking in 2022. The improved techniques for reconstructing the displaced signature of the long-lived HNL are applied to 140.1 fb^{-1} of data used in this analysis collected at $\sqrt{s} = 13$ TeV in 2015–2018. The search results are interpreted in scenarios where the HNL mixes with one more Standard Model neutrino.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

- **AOD**: Analysis Object Data
- **ATLAS**: A Toroidal LHC ApparatuS
- **BSM**: Beyond Standard Model
- **CB**: Combined Muon
- **CL**: Confidence Level
- **CMS**: Compact Muon Solenoid
- **CP**: Combined Performance
- **CR**: Control Region
- **CSC**: Cathode Strip Chamber
- **CT**: Calorimeter-tagged Muon
- **DAOD**: Derived Analysis Object Data
- **DV**: Displaced Vertex
- **ECAL**: Electromagnetic Calorimeter
- **EM**: Electromagnetic
- **FSR**: Final State Radiation
- **GSF**: Gaussian Sum Filter
- **HCAL**: Hadronic Calorimeter
- **HF**: Heavy Flavor
- **HF-CR**: Heavy Flavor Control Region

- **HLT**: High Level Trigger
- **HNL**: Heavy Neutral Lepton
- **IBL**: Insertable B-Layer
- **ID**: Inner Detector
- **IH**: Inverted Hierarchy
- **IP**: Interaction Point
- **ISR**: Initial State Radiation
- **LAr**: Liquid Argon
- **LHC**: Large Hadron Collider
- **LINAC**: Linear Accelerator
- **LLP**: Long-Lived Particles
- **LO**: Leading Order
- **LRT**: Large-Radius Tracking
- **MC**: Monte Carlo
- **MDT**: Monitored Drift Tube
- **ME**: Matrix Element
- **MIP**: Minimum Ionizing Particle
- **MS**: Muon System
- **NH**: Normal Hierarchy
- **NLO**: Next-to-Leading Order

- **pp**: proton–proton
- **PS**: Parton Shower
- **PV**: Primary Vertex
- **QCD**: Quantum Chromodynamics
- **RC**: Random Crossings
- **RF**: Radio Frequency
- **RPC**: Resistive Plate Chamber
- **SCT**: Semiconductor Tracker
- **SFOS**: Same-Flavor Opposite-Sign
- **SM**: Standard Model
- **SR**: Signal Region
- **ST**: Segment-tagged Muon
- **SV**: Secondary Vertex
- **TDAQ**: Trigger and Data Acquisition
- **TGC**: Thin Gap Chamber
- **TRT**: Transition Radiation Tracker
- **VSI**: Vertex Secondary Inclusive
- **WP**: Working Point

INTRODUCTION

"I can thank the fact that I was born at just the right time. A few years older or younger, I would have missed the opportunity... One might call the period from 1925 onward for a few years the Golden Age of Physics when our basic ideas were developing very rapidly and there was plenty of work for everyone to do."

-Paul Dirac, J. Robert Oppenheimer Memorial Prize Speech, 1969

The Standard Model of Particle Physics, in its current form, is the most (or even the only) widely accepted complete theory of the particle nature of the universe. It provides a mathematical description of the three out of four fundamental forces of nature: Weak force, Strong force, and the Electromagnetic force. After its full formulation circa 1975, the model has proved to be extremely successful in predicting natural phenomena at the smallest of scales and highest of energies. Nevertheless, a few but critical open problems challenge the foundations and completeness of the Standard Model, and beg for theoretical extensions that can provide solutions to these problems. A large number of extensions to the Standard Model predicting long-lived particles [1, 2] offer such solutions. At the Large Hadron Collider, long-lived particles can decay far from the proton–proton interaction vertex and produce signatures that are distinct from promptly decaying particles, which constitute most of the searches at the LHC. One such signature is a displaced vertex coming from a long-lived particle decaying within the detector volume of ATLAS producing reconstructible particles. In 2022, the ATLAS experiment's reconstruction algorithms went through a significant overhaul [3, 4] adding unprecedented sensitivity to long-lived particles searches, especially those probing displaced vertices. This improvement in reconstruction, at a time where theories predicting long-lived particles are ubiquitous and relatively unexplored, provides a golden opportunity to search for them using existing colliders, detectors, and dataset.

This thesis presents a new search for long-lived Heavy Neutral Leptons in events with di-lepton displaced vertices taking advantage of the aforementioned improvements in the ATLAS reconstruction algorithm applied to data collected between 2015 and 2018. Chapter 1 outlines the fundamental structure of the Standard Model and lists the open problems that plague it. An extension to the Standard Model in the form of long-lived sterile neutrinos is presented. This extension forms the baselines model used in the analysis described in this dissertation. The apparatus used to test this model is the ATLAS experiment at the Large Hadron Collider, both of which are described in chapter 2, along with the techniques used to collect data and stitch the readout signals from the detector into individual physical objects such as particles and their decays. The overview of this analysis is presented in chapter 3 with the exact goals set out to be achieved. The various steps of the analysis method and the statistical structure of extracting the strength of the model are detailed in chapter 4, followed by the results in chapter 5. Chapter 6 summarizes the inferences of this search and lists out possible improvements in the future versions of this adventure. My contributions to the analysis and to other projects in ATLAS are summarized in chapter A of the appendix.

Chapter 1

THEORY

This chapter describes the theoretical framework which the analysis presented in the dissertation is based on. The Standard Model (SM) is briefly described in section 1.1 along with its limitations in the form of inability to explain observed anomalies. Section 1.2 introduces a family of Beyond Standard Model (BSM) particles, Heavy Neutral Leptons (HNLs), proposed to explain some of these anomalies. HNLs form the benchmark model probed in this search.

1.1 Standard Model

The Standard Model is our best understanding of the particulate nature of the universe at the most fundamental level. The theory[1] considers of two kinds of particles: *fermions*, which have half-integer spins and hence follow Fermi-Dirac statistics, and *bosons*, which have integer spins and follow Bose-Einstein statistics.

Fermions make up all the visible matter while gauge bosons act as the mediators of forces of interaction between them.[2]. The strong force, which is responsible for binding nuclei together, is described by the theoretical framework of Quantum Chromodynamics (QCD) whose mediators are eight *gluons*. The weak force causes radioactive decays to occur mediated by the $W^{\pm}$ and Z bosons. Lastly, the electromagnetic (EM) interaction is mediated by *photons*. The electromagnetic and weak interactions are explained together by the Electroweak (EW) theory, and are commonly known as the electroweak interaction.

The particles of the SM are summarized in fig. 1.1 along with their mass and electric charge. The fermions are organized in three generations (two particles per generation) and two families. The first family is quarks, which have non-zero electric and color charges. Since these are quantum numbers for the EM and QCD interactions, quarks participate in both.

[1] The description of the Standard Model is adapted from the book "Quarks and Leptons" [5], with modified notations for simplicity and self-consistency.

[2] except gravity, for which there is no known particle mediator, and the SM gives no description of the nature of gravity.

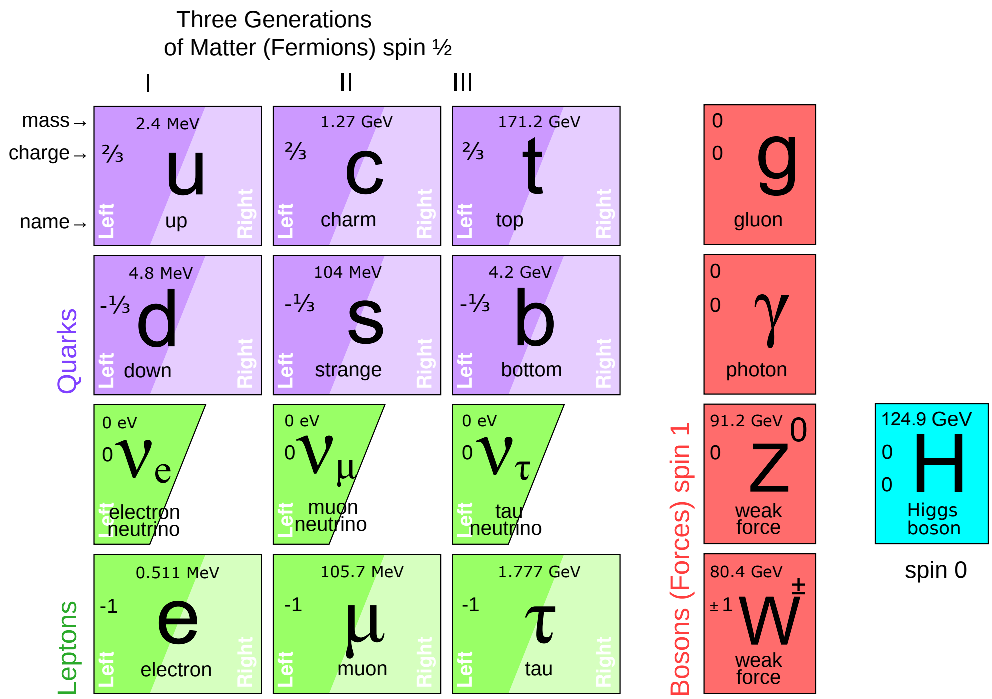

Figure 1.1: The elementary particles in the Standard Model with their mass and electric charge. [6]

The three generations of quarks are *up/down*, *charm/strange*, and *top/bottom*, commonly identified as quark flavors where the flavor is designated by the first letter of their names. The second family of fermions is leptons. Each generation of leptons consists a (electrically) charged-neutral pair, where the charged particles participate in the EM interaction while none of them participate in QCD. The charged leptons are *electrons* (e), *muons* (μ), and *taus* (τ) each having a neutral partner *neutrino* (ν_e, ν_μ, ν_τ). Each fermion has an anti-fermion partner with an indentical mass and magnitude of spin but opposite physical charges. All fermions have non-zero weak isospin quantum numbers (with the caveat of chirality, which is discussed later) and hence participate in the weak interaction. Due to color confinement[3], isolated quarks are not found in nature and instead form color-neutral bound states, called hadrons. Hadrons are typically bound states of two or three quarks (although unstable higher

[3] Gluons themselves have color charges (unlike photons and weak bosons which are not colored in their interactions), which makes the strong interaction a long range force and makes color neutral composite quark states energetically feasible. Increasing the separation between quarks allows for the creation of $q\bar{q}$ pairs, which are also color neutral.

quark multiplicity bound states have also been observed [7, 8]). Hadrons with two quarks (quark-anti-quark pair) are called mesons, for example pions (π^+ is a $u\bar{d}$ bound state), and hadrons with three quarks are called baryons, for example protons (uud) and neutrons (udd).

The last and the most recently discovered piece of the SM is the scalar Higgs boson (H). The Higgs field gives mass to all massive fermions and gauge bosons.

1.1.1 Mathematical Structure

The Standard Model is a mass dimension-4 relativistic quantum field theory with nineteen free parameters. It is defined by the local $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge symmetry. (Roughly,) the three factors of the gauge symmetry give rise to the three fundamental interactions. The first of these is the $SU(3)_C$ gauge symmetry of QCD, where the C stands for color, the quantum number of QCD. The $SU(2)_L \times U(1)_Y$ is the gauge symmetry followed in the EW theory, where the L stands for left-handed, and Y denotes the weak hypercharge, which related to the electric charge Q and the third component of isospin I_3 as $Y = Q - I_3$, the quantum numbers of electroweak interaction.

The quantum states of the fermionic fields are denoted by ψ, which are bi-spinors. They obey the Dirac equation ($(i\gamma^\mu \partial_\mu - m)\psi = 0$) in the Clifford algebra defined using a representation of four 4×4 matrices, called gamma (γ) matrices, which follow the anti-commutation relation

$$\{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu\nu} I_4, \tag{1.1}$$

where $\eta_{\mu\nu}$ is the Minkowski metric and I_4 is the 4×4 unit matrix. Handedness refers to the chiral states of fermions extracted using the projection operator:

$$\psi_L = \frac{1}{2}(1 - \gamma^5)\psi, \;\; \psi_R = \frac{1}{2}(1 + \gamma^5)\psi. \tag{1.2}$$

$\psi_{L(R)}$ is called the left-(right-)handed component of ψ, and $\gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3$ is the fifth gamma matrix. Chirality is a Lorentz-invariant quantity but not a constant of motion, meaning that a pure left-handed state can gain (only for massive fermions) a right-handed component during flight, and vice versa. The left-handed fermions of the same generation form a doublet of the $SU(2)_L$ symmetry. The quark doublets are represented as $Q_L \equiv (u_L \; d_L)^T$ and the lepton doublets are represented as $L_L \equiv (\nu_{e,L} \; e_L)^T$, with similar definitions for all three generations. The right-handed quarks, such as u_R and d_R, and the right-handed charged leptons, such as e_R, are singlets of the $SU(2)_L$ symmetry whereas the right-handed neutrinos, ν_R, would not have charges under any interaction. Right-handed neutrinos have not been observed in nature from direct or indirect measurements, and are not part of the SM. Right-handed fermion particles (and left-handed anti-particles) have a third component of weakisospin $I_3 = 0$, and hence do not participate in the weak interaction. Table 1.1 summarizes the electric charge, weak hypercharge, and the third component of the weak isospin for the fermionic fields.

Field	Q	I_3	$Y = Q - I_3$
$Q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$	+2/3 -1/3	+1/2 -1/2	1/6
u_R	+2/3	0	+2/3
d_R	-1/3	0	-1/3
$L_L = \begin{pmatrix} \nu_{e,L} \\ e_L \end{pmatrix}$	0 -1	+1/2 -1/2	-1/2
e_R	-1	0	-1

Table 1.1: The Standard Model fermion fields and their electric charge Q, third component of isospin I_3, and the weak hypercharge Y.

The $SU(3)_C$, $SU(2)_L$, and $U(1)_Y$ gauge fields are denoted by G^a_μ ($a = 1, ..., 8$), W^b_μ (b = 1, 2, 3), and B_μ, respectively. Lastly, the Higgs field Φ is a complex scalar doublet of $SU(2)_L$. The Lagrangian of the Standard Model is given by:

$$\begin{aligned} \mathcal{L}_{\text{SM}} \equiv & \underbrace{-\frac{1}{4}G^a_{\mu\nu}G^{\mu\nu}_a - \frac{1}{4}W^b_{\mu\nu}W^{\mu\nu}_b - \frac{1}{4}B_{\mu\nu}B^{\mu\nu}}_{\text{gauge term}} + \underbrace{i\bar{\psi}\gamma^\mu D_\mu \psi}_{\text{kinetic term}} + \\ & \underbrace{\bar{\psi}_i y_{ij} \psi_j \Phi + h.c.}_{\text{Yukawa term}} + \underbrace{|D_\mu \Phi|^2}_{\text{Higgs-gauge coupling}} - \underbrace{V(\Phi)}_{\text{Higgs potential}} , \end{aligned} \tag{1.3}$$

where $G^a_{\mu\nu}$, $W^b_{\mu\nu}$, and $B_{\mu\nu}$ represent the field strength tensors of strong, weak isospin, and weak hypercharge fields, and *h.c.* stands for the hermitian conjugate of the previous term. For simplicity, the sum over the fermionic fields in not shown, and the Yukawa term is also simplified. The field strength tensor for a field V is given by

$$V^i_{\mu\nu} \equiv \partial_\mu V^i_\nu - \partial_\nu V^i_\mu + g_V f_{ijk} V^j_\mu V^k_\nu, \tag{1.4}$$

where i, j, and k are indices corresponding to the number of fields for a force, g_V is the coupling constant for the field, and f_{ijk} is the structure constant for the symmetry group. The kinetic terms consists of the covariant derivative D_μ, defined as

$$D_\mu \equiv \partial_\mu - ig_s \frac{1}{2} \Sigma_a \lambda^a G^a_\mu - ig \frac{1}{2} \Sigma_b \sigma^b W^b_\mu - ig' Y B_\mu, \tag{1.5}$$

where g_s, g, and g' are the coupling constants of the strong, weak isospin, and weak hypercharge fields, respectively, λ are the Gell-Mann matrices, and σ are the Pauli matrices which form a representation of the respective symmetry groups.

The Higgs potential is given by

$$V(\Phi) \equiv \frac{\lambda}{2} (\Phi^\dagger \Phi - v^2)^2, \tag{1.6}$$

where λ is a free parameter, and v is vacuum expectation value of the Higgs field. The Higgs mechanism [9, 10] is responsible for the spontaneous breaking of the EW symmetry, which ascribes masses to three of the four gauge bosons and leaves the photon massless. At low energies, the Higgs field acquires a non-zero vacuum expectation value ($< \Phi^\dagger \Phi >= v^2/2$) from the minimization of the Higgs potential and orients itself along an (arbitrary) direction in the field space. The direction is normally chosen to minimize degrees of freedom, and the

Higgs doublet becomes

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v+h \end{pmatrix}, \tag{1.7}$$

where h is an excitation of the Higgs scalar field. The Higgs-gauge coupling term in eq. (1.3) then expands to two terms, one giving rise to couplings between h and the electroweak gauge bosons, and the other providing mass terms for the electroweak gauge bosons, spontaneously breaking the symmetry of the $SU(2)$ electroweak interaction. The mass eigenstates are

$$\begin{aligned} W_\mu^\pm &= \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2), \\ Z_\mu &= \cos\theta_W W_\mu^3 - \sin\theta_W B_\mu, \text{ and} \\ A_\mu &= \sin\theta_W W_\mu^3 + \cos\theta_W B_\mu. \end{aligned} \tag{1.8}$$

where $\theta_W = \tan^{-1}(g'/g)$ is the weak mixing angle and A_μ is the mass eigenstate corresponding to the photon. The masses corresponding to these states are

$$\begin{aligned} m_{W^\pm} &= \frac{gv}{2}, \\ m_Z &= \frac{v}{2}\sqrt{g^2 + g'^2}, \text{ and} \\ m_\gamma &= 0. \end{aligned} \tag{1.9}$$

After the EW symmetry is broken, the Yukawa term decomposes into two terms: an interaction term between the fermions and the Higgs scalar, the strength of which is proportional to the fermion mass, and a Dirac mass term for all (massive) fermions, which is of the form

$$\mathcal{L}_{\text{Dirac mass}} = -m(\bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L). \tag{1.10}$$

Since ν_R do not exist within the SM, no such higgs interaction term mass term exists for neutrinos, causing them to be massless.

1.1.2 Limitations of Standard Model

The Standard Model has survived the test of time and has proved to be the most well-tested theory of natural phenomena. It has enjoyed incredible success in making predictions validated by highly precise measurements, such as the 16 parts-per-trillion validation of the CPT invariance of the Standard Model by the BASE experiment [11], less than 1 parts-per-trillion test of the electron magnetic moment [12], the 0.2 parts-per-mil measurement of the W-boson mass [13] and the measurements of Higgs boson properties by the ATLAS [14] and CMS [15] experiments.

However, quite a few open problems with experimental evidences taint this success. One such problem is that of matter-anti-matter asymmetry, which questions the theoretical basis of the large ratio of matter to anti-matter in the universe. The SM provides no mechanism/reason for this asymmetry coming from interactions in the early univerise. Many astronomical measurements, such as that of the rotational curves of spiral galaxies [16], prove the existence of an invisible kind of matter, called *dark matter*, that forms a larger portion of the universe's energy density than the visible matter predicted in the SM. None of the SM particles can act as dark matter, which poses a question about its particulate nature. The observation most relevant to this thesis and the only direct observation of a BSM phenomena is that of neutrino flavor oscillations. Neutrinos of a specific flavor propagating through space can convert into a different flavor. Such oscillations have been observed in solar neutrinos [17], atmospheric atmospheric neutrinos [18], reactor neutrinos [19], as well as in neutrino beams [20]. Neutrino oscillations can be explained by neutrino mass eigenstates. The Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [21] parameterizes the mixing between neutrino flavor and mass eigenstates. In this framework, the probability of observing a different neutrino flavor than what was originally produced is found to be proportional to the difference of square of masses of neutrino mass eigenstates. The oscillations have been studied in detail, and independent measurements of the PMNS matrix elements have been found to be consistent with each other. Since neutrinos are massless in the SM, these observations strongly suggests a BSM effect giving rise to neutrino masses.

1.2 Heavy Neutral Leptons

Right-handed massive neutrinos are proposed in a large number of BSM extensions in an attempt to provide a complete theory of the universe. They are often referred to as *sterile* neutrinos since they do not engage in any of the gauge interactions. Heavy Neutral Leptons are a similar class of particles that typically contain a superposition of left- and right-handed components and are characterized by the square of couplings (or mixing angles) to SM neutrinos ($|U_e|^2$, $|U_\mu|^2$, and $|U_\tau|^2$). Historically, right-handed neutrinos were introduced as parts of larger gauge groups [22] driven by general principles, such as unified theories [23], to introduce a seesaw mechanism for neutrino masses. A type-I seesaw mechanism provided by HNLs can also explain neutrino masses at the eV scale [24]. The simplest renormalizable extension to the SM consistent with neutrino oscillation measurements adds $\mathcal{N}(\geq 2)$ right-handed $SU(2)_L \times U(1)_Y$ singlet neutrinos. Such minimal extensions are not only adept at explaining neutrino masses, but also offer solutions to other open problems in particle physics such as the baryon-asymmetry in the universe, which can be explained using $\mathcal{N} = 2$ models via leptogenesis [25]. For example, for HNL Majorana masses of the order of the electroweak scale or below, the seesaw mechanism imposes the HNL Yukawa couplings to be small to make the active neutrino masses consistent with flavor oscillation results. The production of HNLs followed by their charge parity-violating decays in the early universe can explain this matter-antimatter asymmetry. [26, 27]. The addition of a light HNL with a very small Yukawa coupling can also provide a dark matter candidate [28, 26]. The rather simple structure of HNL-incorporating models and their versatile use to explain BSM phenomena make their searches very attractive.

1.2.1 Standard Model Extension

The description of the mechanism of generation of neutrino masses via a neutrino portal has been adapted from [29]. The Lagrangian for a minimal extension of the SM, $\mathcal{L}_{\text{SM+HNL}}$, by $\mathcal{N}$ right-handed HNLs N_I is given by:

$$\mathcal{L}_{\text{SM+HNL}} - \mathcal{L}_{\text{SM}} = \mathcal{L}' = i\bar{N}_I \gamma^\mu \partial_\mu N_I - \left(F_{\alpha I} \bar{L}_{L\alpha} N_I \tilde{\Phi} + \frac{M_I}{2} \bar{N}_I^C N_I + h.c. \right), \tag{1.11}$$

where $I = 1, ..., \mathcal{N}$. α is the lepton flavor index running over $\{e, \mu, \tau\}$; $F_{\alpha I}$ is the Yukawa coupling between the α left-handed doublet and N_I; L_L is the left-handed lepton doublet $(\nu_L,\ e_L)^T$; $\tilde{\Phi} = i\sigma_2 \Phi$; M_I is the mass of N_I; and N_I^C is the charge conjugated state of N_I. The extension consists of a kinetic term, the Yukawa term, and a Majorana mass term, $-\frac{M_I}{2}\bar{N}_I^C N_I$, which is allowed since N_I are not charged under any gauge transformations. The SM extension for $\mathcal{N} = 3$ is illustrated fig. 1.2.

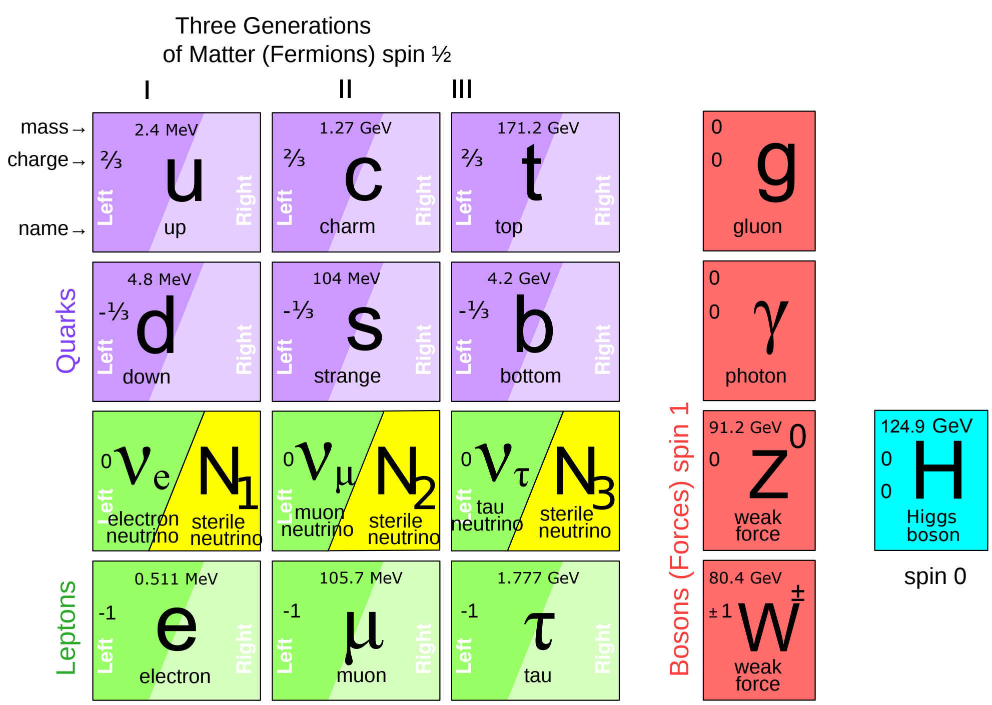

Figure 1.2: The elementary particles in the Standard Model along with $\mathcal{N} = 3$ right-handed Majorana neutrinos. [6]

At low energies, the Higgs field acquires a non-zero v as shown in eq. (1.7), and eq. (1.11) can be written as:

$$\mathcal{L}' = \mathcal{L}_{\text{kinetic}} + \mathcal{L}_{\text{Majorana}} - vF_{\alpha I}\bar{\nu}_{L\alpha}N_I - \frac{F_{\alpha I}}{\sqrt{2}}\bar{\nu}_{L\alpha}N_I h. \tag{1.12}$$

Thus, a Dirac mass term $-vF_{\alpha I}\bar{\nu}_{L\alpha}N_I$ is generated which allows for $\nu_\alpha - N_I$ mixing. The charge eigenstates of eq. (1.11) do not coincide with the mass eigenstates, where the latter can be obtained by diagonalizing the mass-matrix. In the limit $vF_{\alpha I} \ll M_I$, the active neutrino masses (3 eigenvalues) become much smaller than M_I (other $\mathcal{N}$ eigenvalues, up to a small admixture of ν_α) and the electroweak scale v. The smallness of the left-handed admixture in the mass eigenstates is characterized by the ratio of the Dirac and Majorana masses,

$$|\Theta_{\alpha I}|^2 = \frac{v^2|F_{\alpha I}|}{M_I^2} \ll 1, \tag{1.13}$$

where $\Theta_{\alpha I}$ is the mixing angle between left-handed neutrino flavor α and right-handed HNL N_I. This is known as the seesaw mechanism, since a larger HNL mass in comparison to the Dirac mass leads to a lower active mass eigenvalue.

The six free parameters of the PMNS matrix[4] and the three active mass eigenvalues add a total of nine extra parameters to the SM. To accommodate these parameters, $\mathcal{N} \geq 2$ HNLs are required in a valid SM extension. The search presented in this thesis considers two approaches to search for HNLs. A model-independent approach is presented where only a single HNL is kinematically accessible with any other HNLs expected to be beyond the mass scales probed by the analysis. Two single-flavor mixing (1SFH) scenarios are considered in this approach, first where the accessible HNL only mixes with ν_μ (1SFH(μ)) and second only with ν_e (1SFH(e)). A model-dependent approach is also presented, where the results are interpreted in a two-quasi degenerate model (2QDH) [30]. This model considers two HNLs with nearly degenerate masses ($M_1 \sim M_2$, $m_{\text{HNL}} = \frac{M_1+M_2}{2}$), where $|M_1 - M_2|$ is smaller than

[4]Typically, the PMNS matrix is defined with four free parameters consisting of three mixing angles and one charge-parity violating phase. Such a parameterization assumes a Dirac neutrino. For Majorana neutrinos, two additional phases are included for a total of six free parameters.

the detector resolution. The four models considered are tabulated in table 1.2.

HNLs in this analysis refer to the states corresponding to large(r) mass eigenvalues, which consist mainly of the right-handed heavy neutrino N_I and a small admixture of left-handed neutrinos allowing it to participate in the weak interaction. The strength of this interaction is governed by the mixing angles. The experimentally relevant observable measured in this analysis is given by:

$$|U_\alpha|^2 \equiv \Sigma_{I=1}^{\mathcal{N}} |\Theta_{\alpha I}|^2, \tag{1.14}$$

where $\alpha = e, \mu, \tau$. $|U_\alpha|$ is the coupling, or the strength of interaction, between an HNL to a SM neutrino of the flavor α. The total square of coupling strength of HNLs with the SM neutrinos is:

$$|U|^2 \equiv \Sigma_\alpha |U_\alpha|^2 = \Sigma_{\alpha,I} |\Theta_{\alpha I}|^2. \tag{1.15}$$

Since the mixing angles $\Theta_{\alpha I}$ are related to the Yukawa couplings $F_{\alpha I}$ as shown in eq. (1.13), the current limits on these Yukawa couplings obtained using neutrino oscillation data can be used to impose constraints on $|U_\alpha|$ as well. NuFIT provides an updated global analysis of neutrino mixing data [31], which was further analysed in [30] to obtain the corresponding limits on relative mixing angles of HNLs with the three different neutrino flavors, $x_\alpha = |U_\alpha|^2/|U|^2$. The neutrino oscillation data offers two possibilities on the ordering of active neutrino masses:

- Normal Hierarchy (NH): $m_1 < m_2 < m_3$ i.e. $\Delta m_{31}^2 > 0$, and
- Inverted Hierarchy (IH): $m_3 < m_1 < m_2$ i.e. $\Delta m_{31}^2 < 0$.

Figure 1.3 shows the allowed combinations of x_α consistent with neutrino oscillation data. The 2QDH model considered in this analysis is interpreted in an IH and a NH mixing scenario. Typically, most experiments report limits on HNL-ν couplings in a single-flavor mixing scenarios which are more simplistic and hence allow for direct comparisons with other results. Hence, two model-independent 1SFH(μ) and 1SFH(e) scenarios are considered as well. Since τ-leptons are unstable and decay rapidly, they are below the acceptance level of

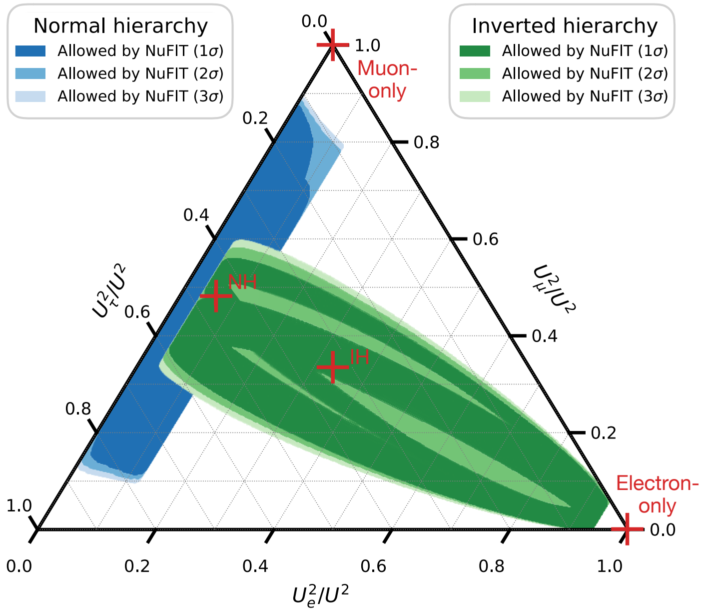

Figure 1.3: Allowed combinations of $x_\alpha = |U_\alpha|^2/|U|^2$ consistent with recent neutrino oscillation data. The combinations allowed by a normal mass ordering is shown in blue and those allowed by an inverted mass ordering are shown in green. The red crosses represent the benchmark scenarios probed in this analysis. Figure adapted from [30].

the analysis, and hence a mixing to them is not considered. The four scenarios probed in this analysis are identified with red crosses in fig. 1.3. Table 1.2 shows the combinations x_α used for the four benchmark models used in this analysis.

Model	x_e	x_μ	x_τ
1SFH(e)	1	0	0
1SFH(μ)	0	1	0
2QDH(IH)	1/3	1/3	1/3
2QDH(NH)	0.06	0.48	0.46

Table 1.2: Relative mixing $x_\alpha = |U_\alpha|^2/|U|^2$ for the four benchmark models used in this analysis.

1.2.2 HNL Production and Decay

This analysis searches for an HNL produced in decays of the W-boson via the $W \to \ell_\alpha N$ mechanism. The HNL production cross-section σ_N depends on the rate of production of the

W boson in proton–proton collisions and the branching ratio of $W \to \ell_\alpha N$:

$$\sigma_N = \sigma(pp \to W) \cdot BR(W \to \ell_\alpha N). \tag{1.16}$$

The branching ratio depends on the mass of the HNL (m_{HNL}), the HNL-ν_α couplings, and the SM branching ratio of $W \to \ell_\alpha \nu_\alpha$, since only the left-handed admixture participates in the weak interaction:

$$\sigma_N = x_\alpha |U|^2 \cdot \sigma(pp \to W) \cdot BR(W \to \ell_\alpha \nu_\alpha) \left(1 - \frac{m^2_{\mathrm{HNL}}}{m^2_W}\right)^2 \left(1 + \frac{m^2_{\mathrm{HNL}}}{m^2_W}\right). \tag{1.17}$$

For the 2QDH model, the production cross-section effectively doubles since the contribution from two HNLs adds up instead of a one single-flavor mixing HNL.

The final state probed in this analysis are the leptonic decays of the HNL ($N \to \ell_\beta \ell_\gamma \nu$), as shown in fig. 1.4. If the charged leptons in the decay are of the same flavor, both charged and neutral currents contribute to the decay. Only decays of W^+ are shown for illustration, but both W^+-and W^--initiated decays are considered in the analysis.

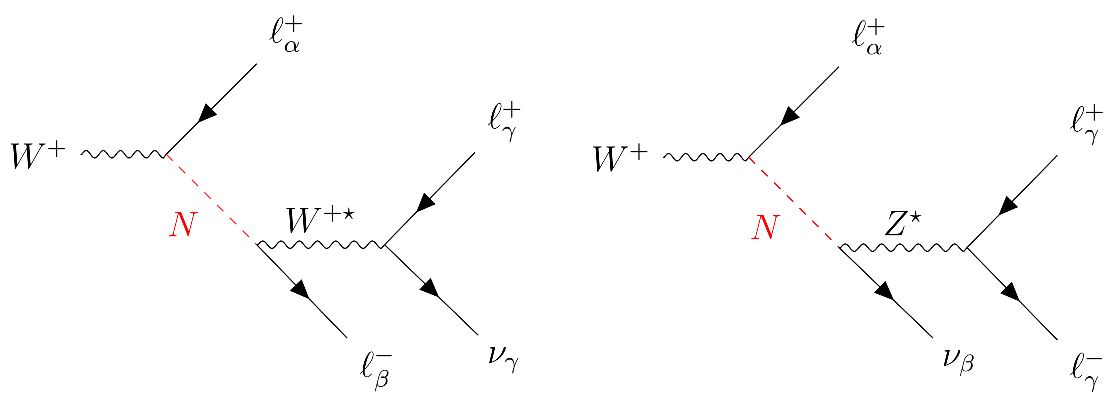

(a) Charged current decay (b) Neutral current decay

Figure 1.4: Feynman diagrams for the W initiated HNL production and subsequent leptonic final state decay.

The SM has an additional approximately conserved quantity during interactions, the total lepton number, given by number of leptons minus the number of anti-leptons. Since the HNLs probed in this analysis have Majorana mass terms in the extended SM Lagrangian, the particle and anti-particle states are identical. Hence, they can give rise to lepton number violating decays as well. Since there are no leptons in the initial state (proton–proton collisions are used), a lepton number conserving decay would lead to a net zero lepton number in the final state as well. Figure 1.5 shows the Feynman diagrams for lepton number conserving and violating decays. An equal contribution from both decays are considered for all models.

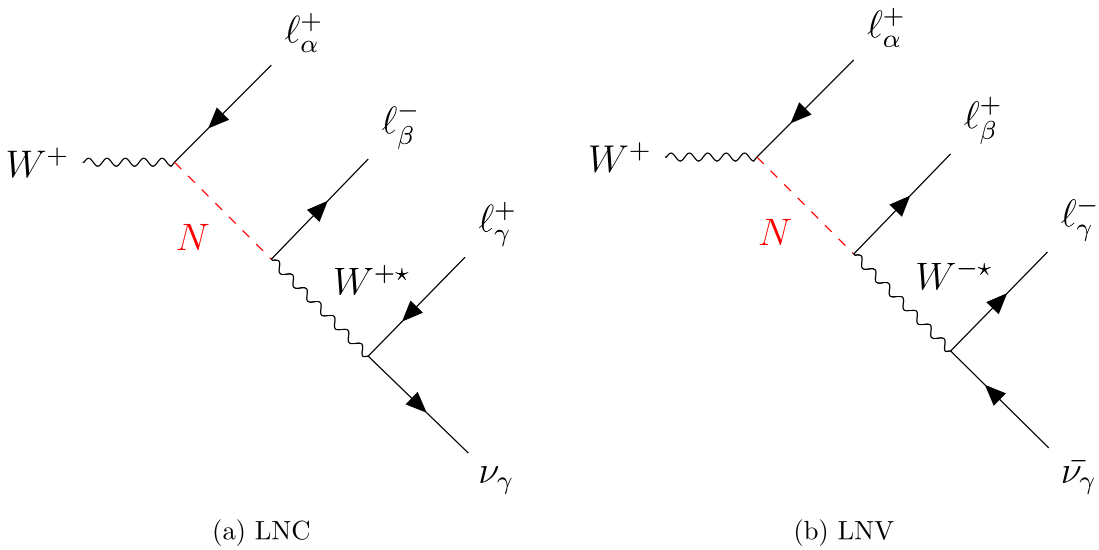

Figure 1.5: Lepton number conserving (LNC) and lepton number violating (LNV) Feynman diagrams for the charged current HNL decay.

The rate of observing an HNL decay in a specific final state is given by the product of its production cross-section and its branching ratio to a final state $BR(N \to \ell\ell\nu)$ where $\ell = e$ or μ. The branching ratio depends on the possible charged and neutral current mediated decays and the coupling strength of the HNL to the different flavors SM neutrinos. The branching ratios to the leptonic final states for the four models considered are shown in fig. 1.6. Contributions from neutrino flavors in the final state are summed for this calculation.

The differences in $BR(N \to \ell_\beta\ell_\gamma\nu)$ and $BR(N \to \ell_\gamma\ell_\gamma\nu)$ are due to interference effects between the charged and neutral current mediated decays. For muon-only mixing HNL,

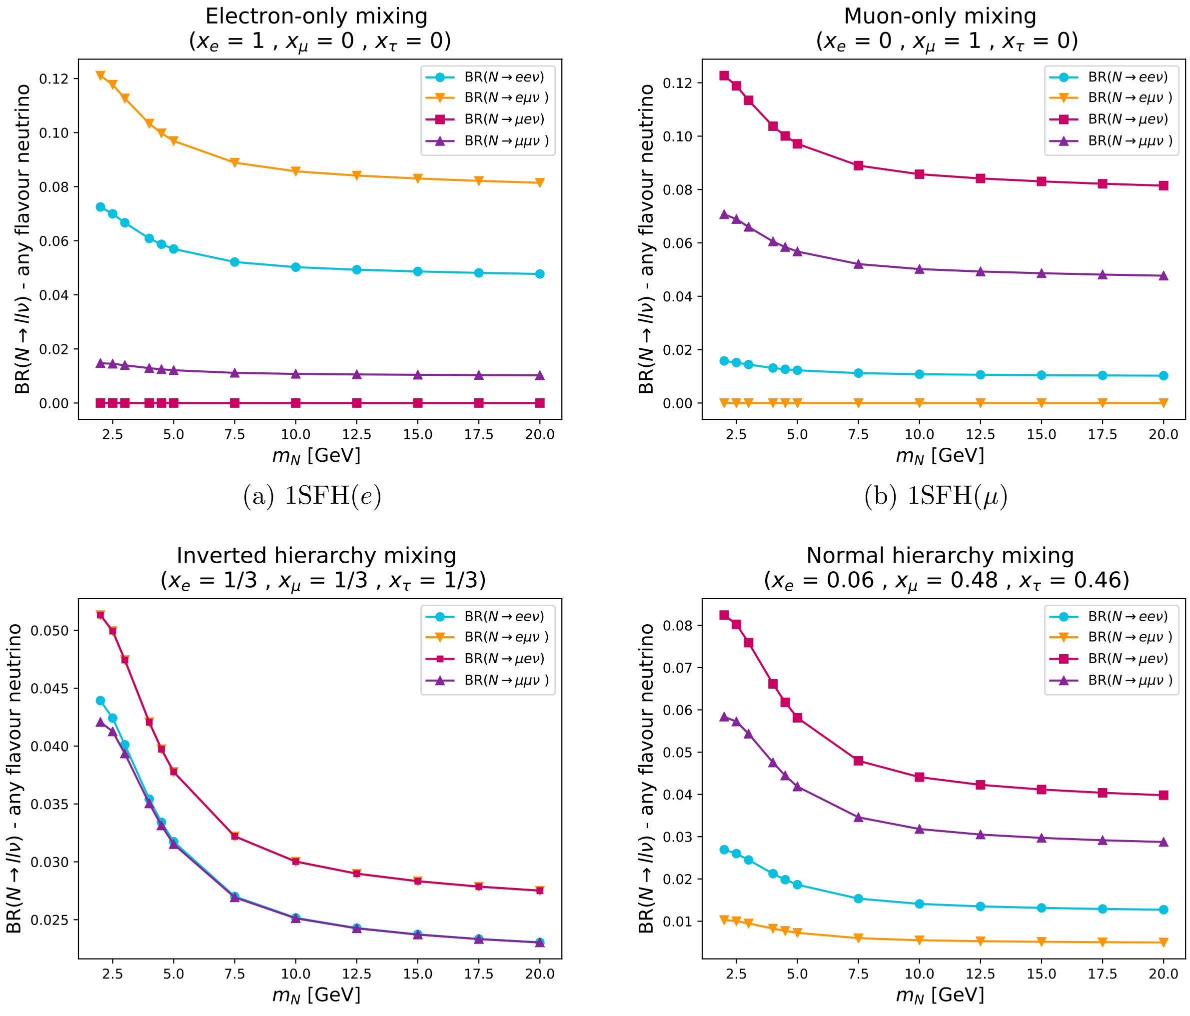

(a) 1SFH(e) (b) 1SFH(μ)

(c) 2QDH(IH) (d) 2QDH(NH)

Figure 1.6: $BR(N \rightarrow \ell\ell\nu)$ for the four benchmark models. For each scenario, the branching ratio is shown for $N \rightarrow ee\nu$ (blue), $N \rightarrow e\mu\nu$ (orange), $N \rightarrow \mu e\nu$ (pink), and $N \rightarrow \mu\mu\nu$ (purple) decays. The y-axis range is adjusted for each model for clarity. [32]

$N \rightarrow e\mu\nu$ is forbidden and $N \rightarrow ee\nu$ is only possible from neutral current decays. Similarly, for electron-only mixing, $N \rightarrow \mu e\nu$ is forbidden and $N \rightarrow \mu\mu\nu$ is possible only from neutral current decays. Since the NH mixing favors a coupling to the muon neutrino, its branching ratios are similar to that of the muon-only mixing scenario, except that all final states have non-zero contributions. The IH mixing favors all final states equally, with the difference coming from the interference between the charged and neutral currents.

The proper lifetime of the HNL, τ_{HNL}, depends on the mass of the HNL, m_{HNL}, and the strength of the coupling to SM neutrinos, U, [33]:

$$\tau_{\mathrm{HNL}} \propto \frac{1}{m_{\mathrm{HNL}}^5 |U|^2}. \tag{1.18}$$

The constant of proportionality depends on the benchmark model being considered. For very small couplings, the lifetime is large enough that it can be resolved by the detectors of the ATLAS experiment. The coupling $|U|^2$ can be analytically computed for a given generated mass and lifetime in a simplified Fermi theory computation [34], approximating the HNL production and decay as a four-fermion interaction vertex. After considering all decay widths and summing over the possible decay channels, the obtained mixing angle, $|U|^2$ as a function of the HNL mass for the muon-only mixing scenario is shown in fig. 1.7 to yield three representative HNL lifetimes. As stated, the shape of the functional relation between $|U|^2$ and m_{HNL} is identical for the models considered, and the proportionality constant has similar magnitudes for them.

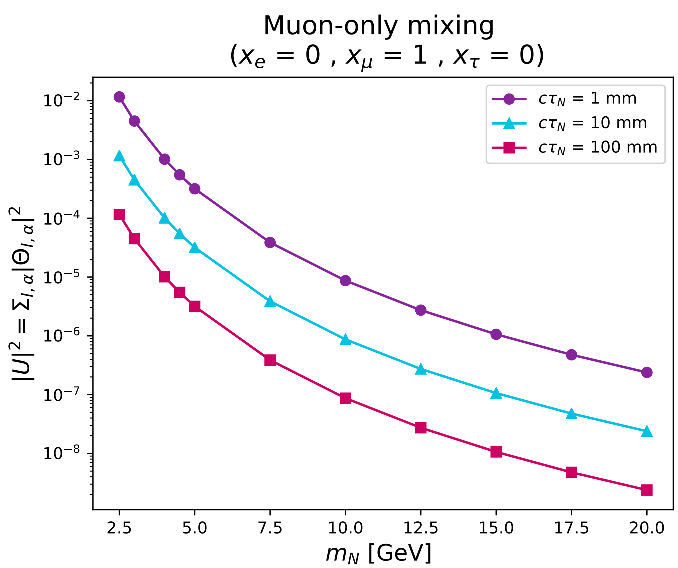

Figure 1.7: The total mixing angle $|U|^2$ as a function of the HNL mass (m_{HNL} or m_N) for the 1SFH(μ) model shown for the $c\tau_{\mathrm{HNL}}$values 1 mm (purple), 10 mm (blue), and 100 mm (pink)). [32]

Chapter 2

EXPERIMENTAL APPARATUS

The results shown in this dissertation are obtained using the data collected by the ATLAS experiment at the LHC, the world's largest particle accelerator hosted by the European Organization for Nuclear Research (CERN). The LHC leads the energy frontier of particle physics colliding protons at unprecedented energies. ATLAS is one of the two multipurpose experiments at the LHC designed to probe a wide range of physics, spanning from precision measurements of the free parameters of the SM to the direct searches of new BSM particles via unique signatures. Its large size along with its sophisticated yet versatile detector systems make it a suitable detector to look for LLPs, including the HNLs probed in this analysis.

The LHC is described in section 2.1 of this chapter. The ATLAS experiment is described in section 2.2 and the reconstruction methods are explained in section 2.3.

2.1 The Large Hadron Collider

The Large Hadron Collider is a two-ring-superconducting hadron accelerator and collider designed to collide protons at a peak center-of-mass energy ($\sqrt{s}$) of 14 TeV and a design luminosity of $L = 10^{34}$ cm^{-2}s^{-1} [35]. It is built as a 26.7 km ring situated at the border of Switzerland and France, and lies between 45 m and 170 m below the surface.

The LHC hosts nine experiments with unique physics goals. The four biggest out of them are LHCb, ALICE, CMS, and ATLAS. The former two experiments target specific phenomena, identifying the origin of matter-anti-matter asymmetry, and understanding the nature of quark-gluon plasma, respectively. The CMS and the ATLAS experiments are general purpose experiments designed to achieve a wide range of physics goals.

Figure 2.1 shows the operation timeline of the LHC and its experiments. Run-1 of the LHC spanned the data-taking period from 2010 to 2012, where the collision energy ranged from 7 TeV to 8 TeV. Run-2 was the data-taking period between 2015 and 2018 at $\sqrt{s} = 13$ TeV, providing the dataset used to perform the analysis described in this thesis. As of April

Figure 2.1: The operation and upgrade timeline of the LHC as of April 2024. [36]

2024, the LHC is in the third year of its third run of data-taking, at 13.6 TeV. Starting 2026, the LHC and the experiments on it will go through a series of upgrades to prepare for the High-Luminosity era of the LHC [37], during which an unprecedented amount (3000 fb^{-1}) of collision data is expected to be collected between 2029 and 2040 by the experiments, pushing the intensity frontier of physics.

The acceleration of protons in the LHC happens in multiple steps, as shown in fig. 2.2. Protons are accelerated using the LINAC-4 (LINAC-2 during Run-1 and Run-2) machine to 50 MeV. The beam is then injected into the Proton Synchrotron Booster, followed by the Proton Synchrotron, which further accelerates the protons to 25 GeV. The pre-final stage is the injection into the Super Proton Synchrotron, where they are accelerated to 450 GeV, before being injected into the two main LHC rings. The beam in one pipe circulates clockwise while the beam in the other pipe circulates anticlockwise. Superconducting dipole magnets are used to create magnetic fields with a strength of about 8 T to further accelerate the protons to up to 6.8 TeV energies.

The hadrons inside the LHC are accelerated using a radio frequency (RF) acceleration scheme [35]. The magnetic field inside the RF cavity osciallates, which affects only the

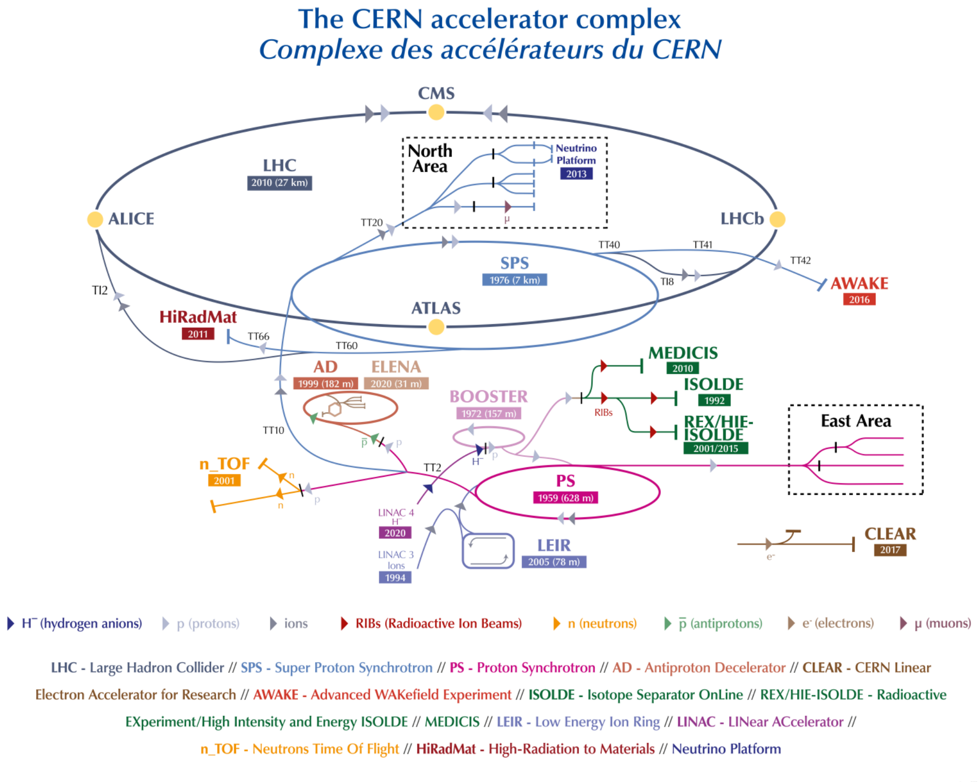

Figure 2.2: The CERN accelerator complex as of 2022.[38]

protons with slightly different energies than the design energy arriving slightly early or late, causing them to either accelerate or decelerate. A direct consequence of this periodic acceleration is the bunch structure of beams. Under nominal operating conditions, each proton beam has 2808 bunches, with each bunch containing about 10^{11} protons. At full luminosity, the LHC uses a 25 ns bunch spacing corresponding to 40 MHz frequency. At the collision point at each of the 4 large experiments, the bunches are squeezed to about 20 μm sizes in the direction transverse to the beam to allow for greater chances of a proton–proton (pp) collision. The number of events N of a specific process produced at the LHC is related to the production the production cross-section of this process σ and the instantaneous luminosity (L) of the beam as:

$$N = \sigma \int L\mathrm{d}t = \sigma \mathcal{L}, \tag{2.1}$$

where $\mathcal{L}$ is the integrated luminosity i.e. L integrated over the data-taking period. The instantaneous luminosity is directly proportional to the square of number of protons in a bunch and inversely proportional to the radius of the bunch. When dense bunches collide at the interaction point (IP), more than one pp collisions occur. On average, the primary collision of interest, the *hard scatter* is the inelastic collision with the most momentum transfer. The other inelastic collisions are called pileup and form a background for a measurement. Figure 2.3 shows the mean number of inelastic pp collisions per bunch crossing for the four years of ATLAS data used in this analysis.

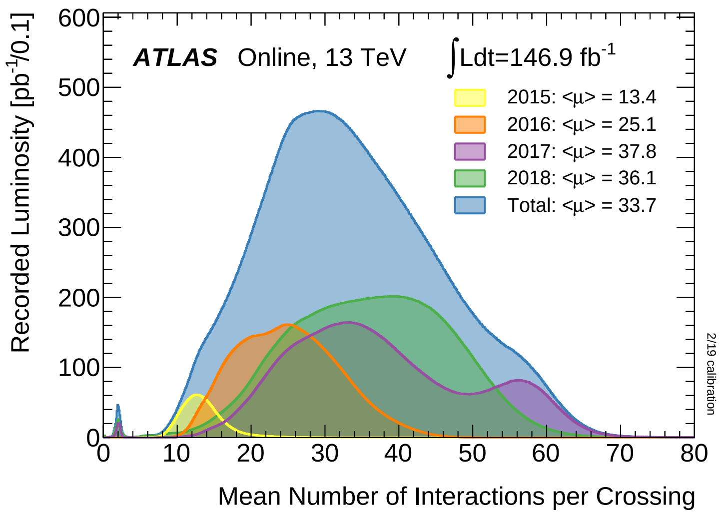

Figure 2.3: Mean number of pp interactions per bunch crossing for each year of data taking in Run-2 and their sum. [39]

2.2 The ATLAS Experiment

The "A Toroidal LHC ApparatuS" (ATLAS) experiment is one of the two general purpose experiments at the LHC, along with the CMS experiment. Situated at Point 1 of the LHC, ATLAS is designed to probe both pp and ion–ion collisions. It is one of the largest detector system in the world with a length of 46 m, and a height and width of 25 m [40]. ATLAS is a cylindrically shaped experiment, built out of concentric detector sub-systems which provide a forward-backward symmetric coverage. Collisions happen at the center of this cylinder producing so-called *outgoing particles*.

The sub-system closest to the IP is the Inner Detector (ID). The ID is responsible for reconstructing the trajectories of charged particles. These trajectories are in-turn used to identify the hard scatter, known as the primary vertex (PV). It is also used to reconstruct secondary decays of non-prompt particles, or secondary vertices (SVs). The ID is surrounded by the calorimeter system, which is used to measure the energy of all particles and stop them, except muons. The outermost detector layer of ATLAS is the Muon Spectrometer (MS) that identifies muons and also provides a secondary momentum measurement. Neutrinos do not interact with any of the detector layers and hence escape ATLAS.Figure 2.4 shows the layout of the ATLAS detector during Run-2.

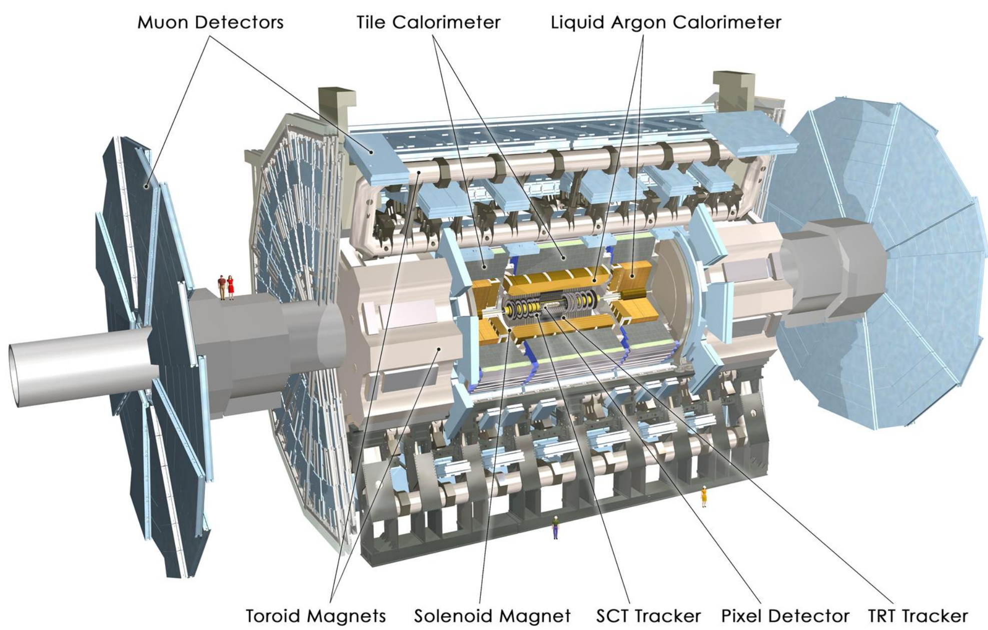

Figure 2.4: Layout of the ATLAS detector during the Run-2 data taking period. [41]

ATLAS features a unique system of four large superconducting magnets. The ATLAS magnet system consists of:

- a solenoid, which is aligned on the beam axis and provides a 2T axial magnetic field for the ID system;
- a *central* toroid, and two *forward* toroids, which produce a toroidal magnetic field of 0.5T and 1T for the muon detectors in the central and forward regions, respectively.

The exact magnetic field strengths and directions are measured using hall probes and the magnetic field map build out of this measurement is used for reconstruction.

2.2.1 ATLAS Geometry

ATLAS follows a right-handed coordinate system to define the kinematics of objects in its measurements. The beamline defines the longitudinal z-axis, with the xy-plane transverse to it. The x-axis points towards the center of the LHC, and the y-axis points towards the Earth's surface. Figure 2.5 illustrates the coordinate system used in ATLAS.

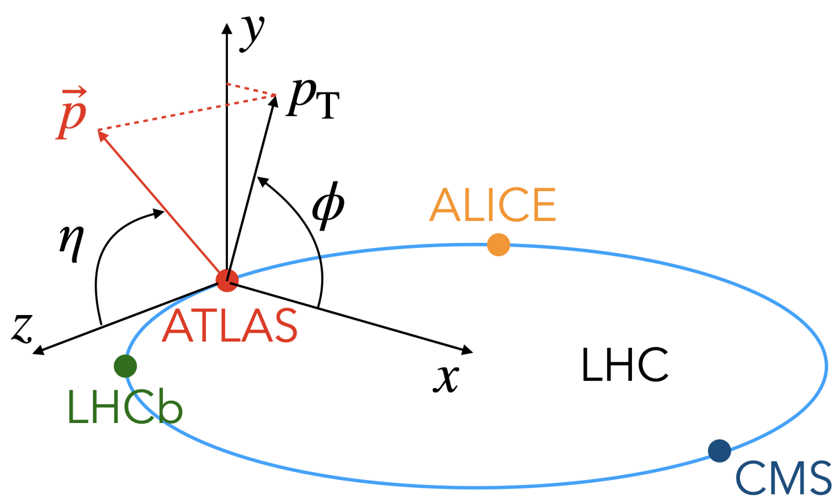

Figure 2.5: Geometry followed by ATLAS to describe directions and kinematics of particles.

The angle in the transverse plane with respect to the x-axis is the azimuthal angle (ϕ). The angle w.r.t. the beam direction is the polar angle (θ), and is more commonly represented by the pseudo-rapidity (η) defined as:

$$\eta = -\ln(\tan(\theta/2)), \tag{2.2}$$

Angular distance is measured in units of $\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2}$[1]. The projection of the

[1] The exact definition is $\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta y)^2}$, where y is the rapidity. At relativistic energies, $y \approx \eta$ and hence η is generally used instead y.

momentum in the transverse plane, the transverse momentum (p_{T}), is defined as:

$$p_{\mathrm{T}} = \sqrt{(p_x)^2 + (p_y)^2}. \tag{2.3}$$

The ATLAS detector has a nearly full coverage in ϕ, and a coverage in pseudo-rapidity up to $|\eta|$=4.9. The wall of the cylinder is also referred to as the **barrel**, and the side faces as the **endcaps**. The exact η transition from barrel to endcap depends on the detector subsystem and shall be defined in relevant contexts throughout this chapter.

2.2.2 Inner Detector

The Inner Detector is the sub-system closest to the collision point. Figure 2.6 shows the cross-section of the ID system in the barrel.

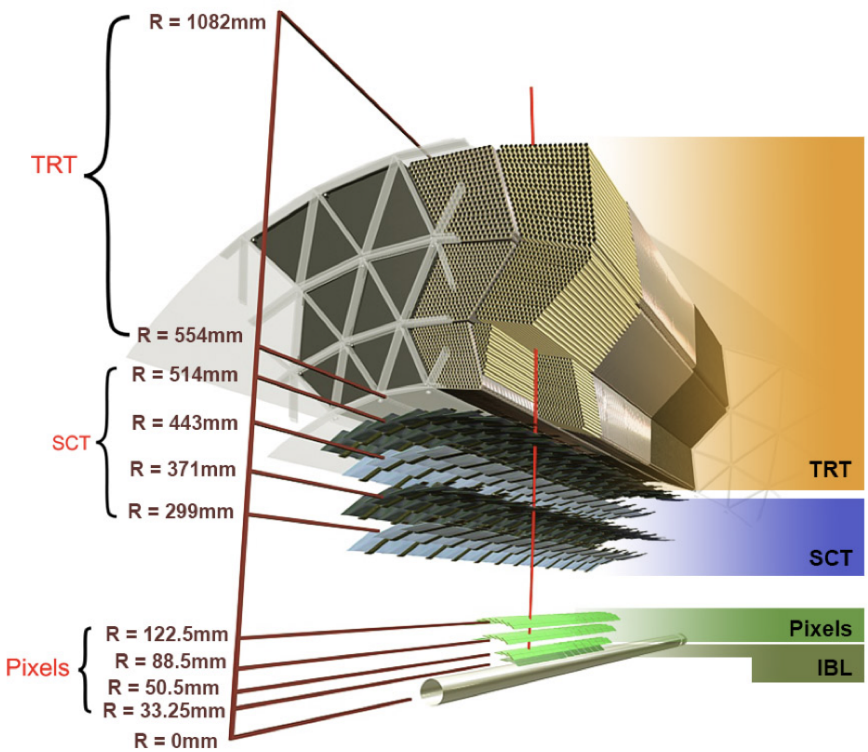

Figure 2.6: Illustrative cross-section of the Inner Detector in the barrel region. [3]

The innermost layer is a pixel silicon detector, the Insertable B-Layer (IBL) [42], which was installed between Run-1 and Run-2 to help retain high reconstruction efficiency in high density pp collision environments. The IBL is surrounded by three additional layers of pixel detectors. These four layers form the Pixel System. Having high-granularity detectors close to the collision point ensures efficient identification of the PV. The four layers of silicon micro-strip modules form the Semi-Conductor Tracker (SCT) system with a total of 4088 modules. Each barrel or disk provides two strip measurements from back-to-back modules glued at a stereo angle (40 mrad) which are combined to build a space-point. The outermost layer of the ID is the Transition Radiation Tracker (TRT) made out of approximately 300,000 thin-walled drift/straw tubes made out of kapton, each having a a gold-plated Tungsten wire 31 μm in diameter [43]. The TRT is filled with an Argon-based gas mixture which is ionized by passing charged particles. Table 2.1 outlines the main characteristics of the different ID detectors in the barrel., The ID has a coverage roughly up to $|\eta|$=2.5 from the silicon detector layers, and up to $|\eta|$=2.1 from the gaseous detector.

Sub-detector	Active Element Size $r-\phi\times z$	Resolution [μm]	Hits in the barrel	Radius of the barrel layers [mm]
IBL	50×250 μm	10×66.5	1	33.25
Other Pixel	50×400 μm	10×115	3	50.5, 88.5, 122.5
SCT	80 μm	17	8	299, 371, 443, 514
TRT	4 mm	130	∼30	554-1082

Table 2.1: Main characteristics of the four ATLAS ID sub-detectors. [44, 45]

2.2.3 Calorimeters

ATLAS consists of two calorimeter systems, the Electromagnetic Calorimeter (ECAL), and the Hadronic Calorimeter (HCAL), which are designed to detect electromagnetic and hadronic showers, respectively. When high energy electrons interact with dense material, they radiate photons via bremsstrahlung [46], and the photons in-turn produce electron pairs forming a shower. The shape and depth of the resulting shower depends on the incident energy of the electron, the resolution of which depends on the granularity of the calorimeter cells with added effects from the energy deposit and detector noise. The ECAL is designed to have a high granularity in the η range of the ID for maximum photon and electron identification efficiency. Hadronic showers are, instead, intiated by a collision of an incident hadron against a nucleus. Such a collision leads to a decay with many hadrons and an excited nuclear state, which in turn decays decays into lighter hadrons. This cascade leads to a shower of hadrons which can interact via both, nuclear and electromagnetic reactions. Hence, the resulting showers are then measured in combination by the ECAL and HCAL. The ATLAS calorimeters can measure the kinetic energies of all particles that reach them, except for neutrinos and muons, since the former do not interact with the calorimeters at all and the latter leave $\sim$3 GeV energy in the calorimeter and pass through. Figure 2.7 shows the layout of the calorimeters in the ATLAS detector.

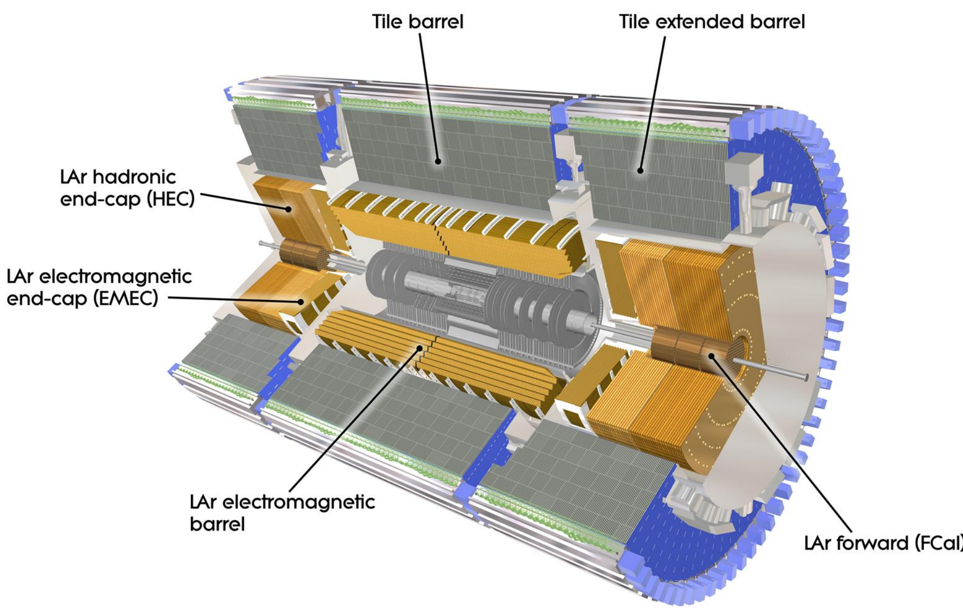

Figure 2.7: Computer generator layout of the ATLAS Calorimeter subsystem. [47]

The ECAL is divided into a barrel part ($|\eta|$ <1.475) and two end-cap components ($1.375 < |\eta| 3.2 <$). It is a hybrid lead-Liquid Argon (LAr) detector with accordion-shaped kapton electrodes and lead absorber plates over its full coverage. This geometry provides a full ϕ coverage. In the region of $|\eta| < 1.8$, a presampler detector, made out of an active LAr layer of thickness 1.1 cm (0.5 cm) in the barrel (endcap) region, is used to correct for the energy loss of the electrons and photons upstream of the calorimeter.

The HCAL uses different detector technologies in different parts of ATLAS. In the barrel ($|\eta| < 1$) and extended-barrel ($0.8 < |\eta| < 1.7$), steel is used as the absorber material and scintillating tiles as the active material. The Hadronic Endcap Calorimeter lies within $1.5 < |\eta| < 3.2$ and uses an interleavened copper-LAr material. Finally, the LAr Forward calorimeter extends the range to $3.1 < |\eta| < 4.9$, and serves as both an electromagnetic and hadronic calorimeter, with one layer of copper and two layers of tungsten, with LAr between the layers acting as the sensitive medium.

2.2.4 Muon System

At the relativistic momentum values ($\mathcal{O}$(1-100) GeV/c) with which most muons are reconstructed in ATLAS, muons act as minimum-ionizing particles (MIPs) [46]. This causes them to lose only ~3 GeV of energy in the calorimeter, and pass through all ATLAS detector systems. The ATLAS MS is a tracking subsystem providing a secondary momentum measurement (in addition to the measurement from the ID) for muons, and most importantly particle identification, since muons are the only charged particles that are not stopped by the calorimeters.

Figure 2.8 shows the layout of the ATLAS MS. The MS consists of three concentric chamber layers in the barrel region and three disk layers in the endcap region. The toroid magnets provide the bending power in the MS, and the trajectory is reconstructed using signatures left in different sub-detectors.

Over most of the $|\eta|$ range, a precision measurement of the track coordinates in the principal bending direction of the magnetic field is provided by the Monitored Drift Tubes (MDTs). MDTs have a coverage up to $|\eta|$=2.7 (2.0 for the innermost layer), and they use pressurised drift tubes with a diameter of 29.970 mm operating with an Ar/CO_2 mixture to collect muon hits. A central tungsten-rhenium wire with a diameter of 50 μm passes through

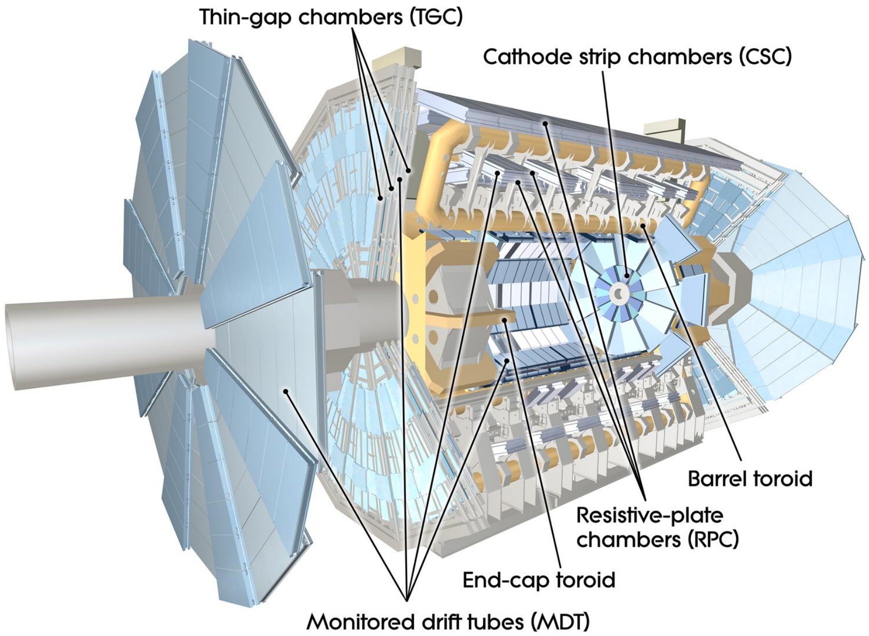

Figure 2.8: Computer generator layout of the ATLAS Muon subsystem. [48]

the tubes. MDTs work on the same principle as the TRT, where drift circles left created by passing charged particles are measured by the electrodes to reconstruct their position. The technology and geometry allows for an average spatial resolution of 80 μm. Due to the relatively higher fluence of particles in the forward (2.0< $|\eta|$ <2.7) region, a second set of detectors, Cathode Strip Chambers (CSCs), are used to maintain muon reconstruction efficiency. The CSC has perpendicularly running cathode and anode wires. A charged particle passing through the CSC knocks electrons off of the gas which flock to the anode and create a signal pulse in the cathode strips, giving a measurement of two position coordinates. The triggering of events (see section 2.2.5) in the MS is facilitated by Resistive Plate Chambers (RPCs) in the barrel and Thin Gap Chambers (TGCs) in the endcap. These detectors also provide a measurement of the second coordinate, which is the ϕ angle. Table 2.2 summarizes the coverage and functions of the ATLAS MS sub-detectors.

MS Sub-detector	Coverage	Function
Muon Drift Tubes	$\|\eta\| < 2.7$ (innermost layer: $\|\eta\| < 2.0$)	Precision tracking
Cathode Strip Chambers	$2.0 < \|\eta\| < 2.7$	Precision tracking
Resistive Plate Chambers	$\|\eta\| < 1.05$	Triggering, second coordinate
Thin Gap Chambers	$1.05 < \|\eta\| < 2.7$	Triggering, second coordinate

Table 2.2: η coverage and function of the MS sub-detectors.

2.2.5 Trigger and Data Aquisition System

The LHC delivers a proton bunch every 25 ns, and each bunch is tightly packed with many protons which collide to give a large fluence of outgoing particles. The ATLAS experiment cannot record all of these events because of limitations imposed by the data-taking rate of the detector readout electronics and the storage and processing capacity of the computing infrastructure in place. Hence, a trigger system is used to record only the interesting physics events, where dedicated triggers are designed and optimized based on the targeted physics signatures.

The ATLAS Trigger and Data Acquisition (TDAQ) system is responsible for the so-called *online*[2] processing, selecting and storing of events of interest for *offline*[2] analysis, and is shown diagrammatically in fig. 2.9. The TDAQ system runs in two levels: the hardware-based Level-1 (L1) trigger, and the software-based High Level Trigger (HLT). The L1 trigger runs on hardware computational elements thanks to their high parallelization capacity and processing speed. Latency requirements [49] imposed by the rate of collisions also require that the first level of processing happen within the cavern, since light is not fast enough to carry the data outside the cavern to the surface before the next event occurs. The HLT, on the other hand, can afford slower speeds thanks to the reduced throughput coming out of L1, and allows for the combination of event information from different parts of the detector [49]. In the HLT, a typical reconstruction sequence makes use of dedicated fast trigger algorithms to provide early rejection, followed by more precise and more CPU-intensive algorithms that

[2] Online analysis refers to the (multi-)object reconstruction which takes place during data-taking for the purpose of trigger selection. Offline analysis refers to the full event reconstruction which happens downstream, after a decision has been made to retain the event.

are similar to those used for offline reconstruction to make the final selection [49].

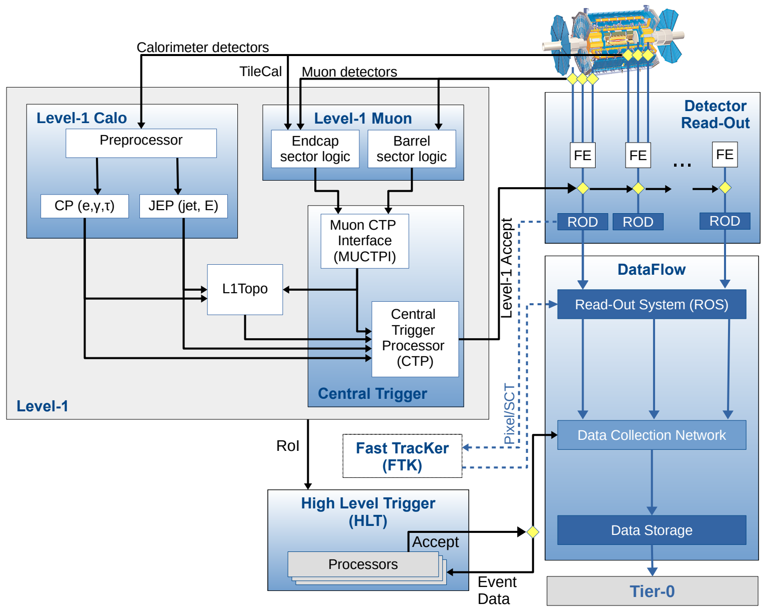

Figure 2.9: Schematic of the ATLAS Trigger and Data Acquisition System in Run-2. (Figure adapted from [49])

The L1 uses information from the calorimeter and muon detectors [49]. The L1 calorimeter trigger takes signals from the calorimeter detectors as input while the L1 muon trigger uses hits from the RPCs (in the barrel) and TGCs (in the end-caps). To reduce the rate in the endcap region of particles not originating from the interaction point, the L1 muon trigger applies coincidence requirements between the outer and the inner TGC stations, as well as between the TGCs and the tile calorimeter [49].

The L1 trigger accepts events at a rate up to the maximum detector-readout rate of 100 kHz, down from the bunch crossing rate of about 40 MHz. The HLT further reduces this physics output rate to an average of 1.2kHz.

Combined with the average size of a physics event, the average physics throughput to data storage comes to 1.2 GB/s [49]. Once an event is accepted by the HLT, the data is sent to permanent storage for offline reconstruction at CERN's computing centres.

2.3 Particle Reconstruction

ATLAS can identify and reconstruct all known stable (except neutrinos) and some unstable particles. Figure 2.10 illustrates the signatures left behind by particles passing through ATLAS. These signatures (such as hits or energy clusters) are *stitched* together to identify the nature of the particle that left them and reconstruct its kinematics. This section discusses the standard reconstruction methods ATLAS uses to build so-called *physics objects*.

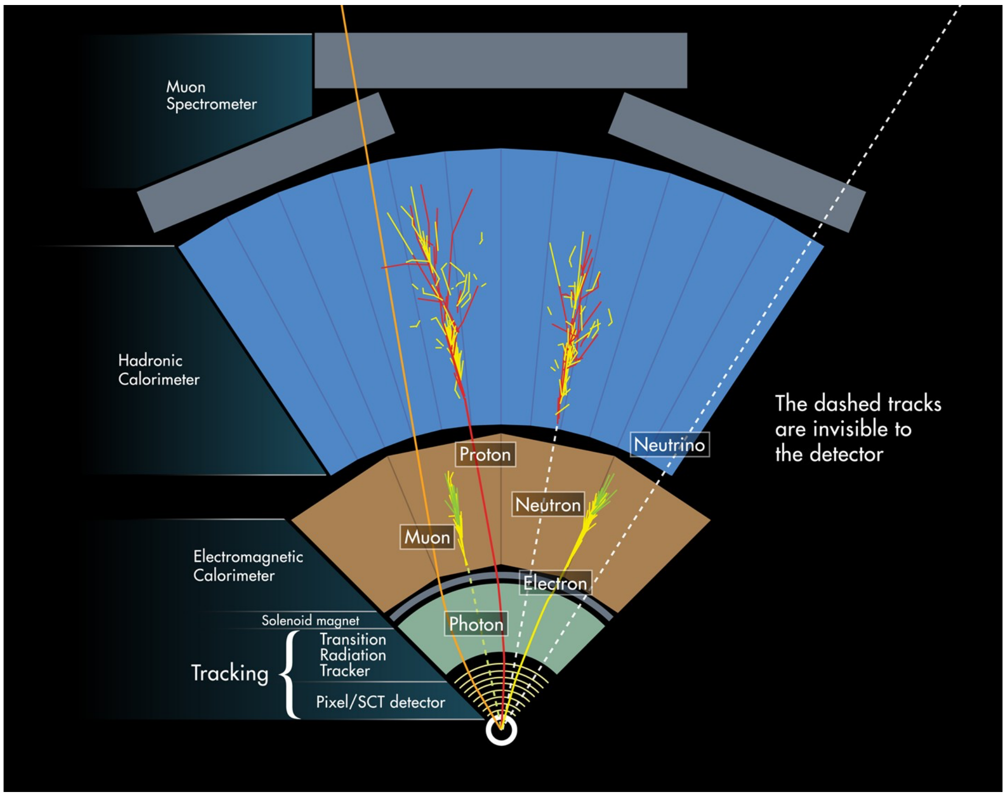

Figure 2.10: Simplied illustration of signatures left by stable particles in ATLAS. [50]

As described in section 2.2, the ATLAS sub-detectors are generally either *trackers* (ID and MS), or calorimeters (ECAL and HCAL). Tracking refers to the identification of the series of hits that have been left behind by the a charged particle followed by the measurement of its trajectory using these hits. Calorimetry refers to the measurement of the direction and energy of a particle using clusters of calorimeter cells. The two methods provide complementary benefits, since the reconstructed momentum resolution from tracking gets worse with increasing particle momentum, whereas the energy resolution from calorimetry improves with increasing particle energy. The design inverse transverse momentum resolution [51] for the ID is

$$\sigma\left(\frac{1}{p_{\mathrm{T}}}\right)\cdot p_{\mathrm{T}} = 0.036\%\cdot p_{\mathrm{T}} \oplus 1.3\%, \tag{2.4}$$

while the calorimeter energy resolution [51] for single charged pions in the center of the detector is

$$\frac{\sigma\left(E\right)}{E} = \frac{50\%}{\sqrt{E}} \oplus 3.4\% \oplus \frac{1\%}{E}, \tag{2.5}$$

where energies and transverse momenta are measured in GeV.

2.3.1 Tracks

Tracks are reconstructed trajectories of charged particles using hits left behind by them in the ID system of ATLAS. Track reconstruction is a critical step in the ATLAS analysis chain. Tracking gives a precise measurement of the momentum and the direction of charged particles, and provides an extrapolation of the trajectory to the IP, where a particle is nominally expected to originate from.

As shown in fig. 2.11, tracks are described by 5 helix parameters in a representation around the beam line (the ATLAS z-axis) with the reference point at the average measured beamspot position. Two of the parameters are the transverse ($\boldsymbol{d_0}$) and longitudinal ($\boldsymbol{z_0}$) impact parameters, defined as the transverse and longitudinal coordinate of the single point of closest approach transverse to the reference point. The other three parameters are the azimuthal ($\boldsymbol{\phi}$) and the polar ($\boldsymbol{\theta}$) angles of the track momentum at the reference point, and the ratio ($\boldsymbol{q/p_{\mathrm{T}}}$) of the charge of the reconstructed track divided by the magnitude of its momentum projected orthogonal to the magnetic field.

Track reconstruction happens in several steps, and in two *passes*. The first pass, or *inside-out* reconstruction, is the primary tracking method where the track finding is initiated in the pixels and SCT and then and TRT hits are added. The secondary pass, or *outside-in* reconstruction, starts from the TRT and expanded to the silicon sub-detectors. The secondary

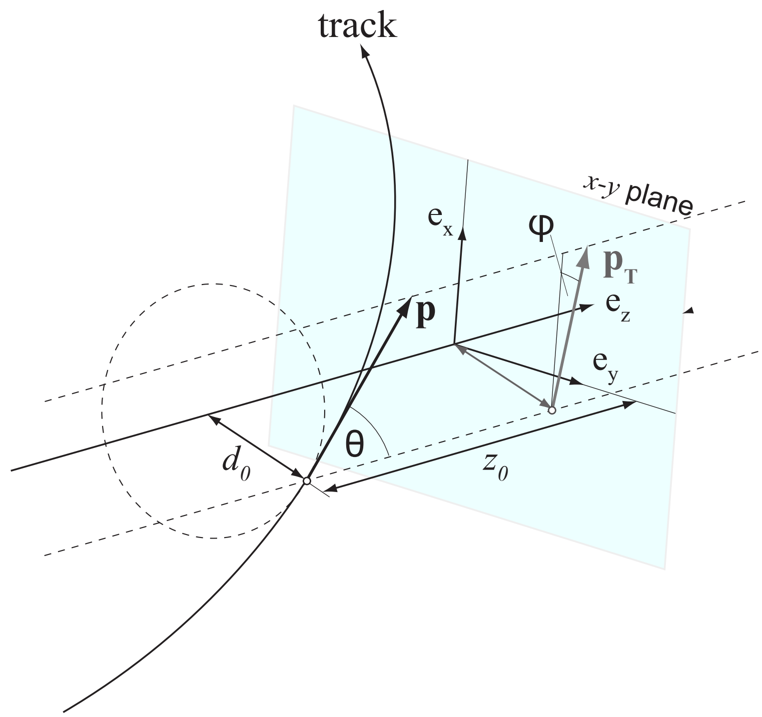

Figure 2.11: The parameters used to represent the helix trajectory of a charged particle with respect to the beamspot reference. [52]

pass aims at reconstructing tracks which start away from the IP using hits left behind after the primary pass. Specifically, it is designed to recover tracks from $\gamma \rightarrow e^+e^-$ conversions, and hence track reconstruction is only attempted in regions of interest determined by deposits in the ECAL. In the first step for both passes, hits from physically adjacent channels in the pixel and SCT sub-detectors are combined to form clusters. These clusters are expected to be left behind by a single charged particle. In the TRT, drift circles are formed from the measured drift times of the gas by a charged particle.

For the primary pass, seeds are formed using triplets of space-points either in the pixel or SCT, and a combinatorial Kalman filter [53] is used to extend these seeds within a narrow search window based on an expected trajectory of the charged particle. An ambiguity resolution step follows which aims at partially resolving shared hits between track candidates, and a global χ^2 minimizing fit is performed to obtain a high-precision track parameter estimate. Finally, an extension to the TRT sub-detector is performed where compatible TRT hits are added before performing a final χ^2 minimizing fit.

Seeding for the secondary pass is triggered by segments of hits in the TRT compatible with the ECAL deposits, and doublets of silicon space-points are constructed for the seeds instead of triplets. A similar Kalman filter based extension followed by ambiguity resolution and χ^2 fit is performed, concluding with an extension to TRT. Figure 2.12 summarizes the steps used in the track reconstruction for both the primary and back-tracking passes in ATLAS.

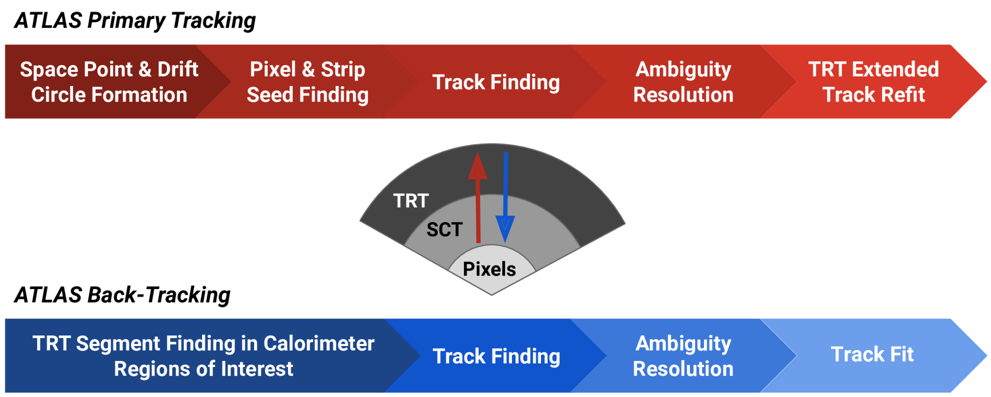

Figure 2.12: Simplified overview of the primary tracking chain and secondary back-tracking chain used in ATLAS ID track reconstruction. The primary reconstruction runs from inside-out, starting from silicon space-points in the innermost Pixel and SCT subdetectors. The secondary back-tracking chain runs from outside-in, seeded from leftover TRT hits within electromagnetic calorimeter regions of interest. [3]

Given the high (and increasing) *pileup* in ATLAS, the precise reconstruction of the multiple interactions and the accurate identification of the collision of interest (generally speaking, there's only one collision of interest) are extremely important for physics analyses. The locations of these interactions are determined using an adaptive multi-vertex fitter [54]. For most physics analyses in ATLAS, the PV is chosen as the vertex with the highest $\boldsymbol{\Sigma p_{\mathrm{T}}^2}$, where the sum is over all the tracks attached to the vertex.

2.3.2 Large-Radius Tracking

The two standard tracking passes in ATLAS, primary and back-tracking, are highly efficient in reconstructing tracks from most physics signatures, except those from LLPs. The primary pass doesn't cover tracks that do not originate from the primary *pp* collisions, and the secondary pass, in addition to pointing requirements, is limited in coverage by the ECAL driven region of interest determination. Hence, there is significant loss in reconstruction efficiency for signatures from LLPs decaying within the ID volume. To recover this efficiency, ATLAS deploys a dedicated large-radius tracking (LRT) pass with loosened IP pointing requirements on the tracks.

The primary tracking pass imposes a strict cut on the track candidate's transverse impact parameter ($|d_0| \leq 5$ mm) to satisfy pointing requirements to the PV. This cut is significantly relaxed ($|d_0| \leq 300$ mm) in the LRT pass. The LRT pass only uses the hits left after running the primary and the secondary passes for track reconstruction which limits possible combinatorics. To counterbalance the effect of such a loosening, other cuts are made stricter which determine the quality of the reconstructed tracks, having a net positive balance between the reconstruction efficiency of tracks from LLPs and the rate of fake combination of hits reconstructed as tracks. Similar to the primary pass, LRT runs only in an *inside-out* fashion, with the seeding step only making use of SCT space-points. Fewer holes[3] are allowed in LRT, and no double holes[4] are allowed. Furthermore, the width of the window used for to search for modules for hits during seed extrapolation (road width) is much narrower in the LRT pass (5 mm vs 12 mm). Table 2.3 summarizes the main differences between the reconstruction criteria of the primary tracking pass and LRT.

Figure 2.13 shows the track reconstruction efficiency of muons or electrons for simulated long-lived Heavy Neutral Lepton (m_{HNL} = 10 GeV, $c\tau_{\mathrm{HNL}}$= 10 mm) production events as a function of the true $|d_0|$ of the HNL decay products, which are two charged leptons (two electrons, two muons, or one electron and one muon) and a neutrino for the study in the plot. As expected, the efficiency from the standard tracking passes (in pink) fall sharply after $|d_0| = 5$ mm, whereas the LRT pass recovers the loss to a very high degree (in blue).

[3] A hole is defined as a missing hit on an active detector module, which was expected to be present between the first and last hits based on the particle trajectory.

[4] A double hole is defined as two consecutive missing hits on active detector modules where both were expected between the first and the last hit based on the particle trajectory.

Selection Criteria	Primary (*inside-out*)	LRT
max. $\|d_0\|$ [mm]	5	300
max. $\|z_0\|$ [mm]	200	300
min. p_{T} [GeV]	0.5	1
max. $\|\eta\|$	2.7	3.0
max. silicon holes	2	1
max. double holes	1	0
max. holes gap	2	1
road width [mm]	12	5
seeding	Pixels and SCT	SCT only
max. seeds per middle Pixel space-point	1	-
max. seeds per middle SCT space-point	5	1

Table 2.3: Most important selection criteria that differ between the inside-out primary tracking and LRT setups. [4]

The slight dip in total efficiency after $|d_0|$ =5 mm is due to the stricter requirements imposed on tracks in the LRT pass, as outlined in table 2.3.

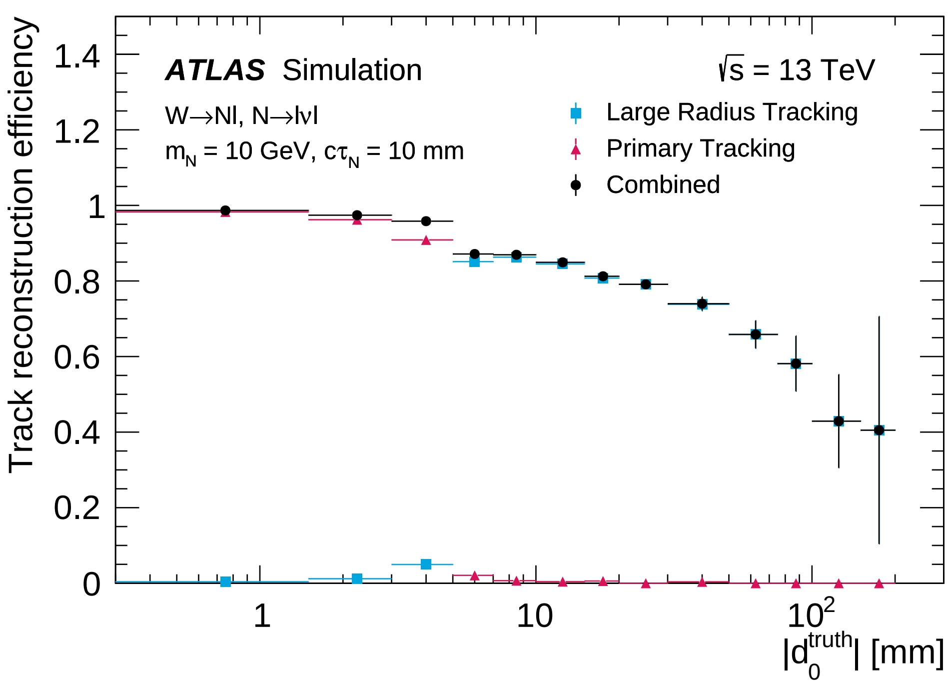

Figure 2.13: Track reconstruction efficiency as a function of the true track $|d_0|$ for simulated Long Lived Heavy Neutral Lepton (m_{HNL} = 10 GeV, $c\tau_{\mathrm{HNL}}$= 10 mm) production and decay events. [4]

2.3.3 Muons

Like track reconstruction, muon reconstruction also happens in several steps and in two passes [55]. The primary reconstruction pass runs *outside-in*, starting from hits in the MS followed by an attempted combination with ID tracks. The secondary reconstruction pass runs *inside-out* attempting to recover muons with low p_{T} which do not leave clean tracks in the MS. The first step common to both passes is cluster formation, where signals from adjacent channels in the CSC are combined. Muons are not expected to leave hits in multiple drift tubes of the same layer. Straight line *muon segments* are formed using a Hough transform in the MS in each of the tracking layers. ID tracks are pre-selected based on minimum silicon hit requirements, and are extrapolated to the inner surface of the MS to obtain their expected 3-momentum at the intersection point.

For the primary pass, segments from all MS layers are combined to form MS tracks. Seed segments are formed in the middle MS layer, and an extrapolation is done based on an assumed parabolic trajectory of the muon. An ambiguity resolution step follows to resolve tracks which share hits, and a global χ^2 minimizing fit is performed to obtain the final MS track. MS tracks with $|\eta| < 2.5$ are combined with ID tracks to improve the track-parameter resolution. ID tracks in a narrow angular window of compatibility wit the MS tracks are selected and combined with the MS track in a final global χ^2 fit which also accounts for the energy loss of the muon in the calorimeter. MS information may be added or removed from the candidate based on this final fit quality. The primary pass yields three different types of muons:

- **Combined (CB) muons** with a full MS track and a full ID track,
- **Silicon-associated Forward (SiF) muons** beyond $|\eta|$ =2.5 which have an MS track associated with a "tracklet" consisting of only pixel hits, but not a full ID track, and
- **Standalone (SA) muons** with a full MS track but no associated ID track or hits.

While the primary pass reconstructs the majority of the muons used in ATLAS, a secondary recovery chain is run to recover low p_{T} muons and muons in regions where the MS lacks coverage. The reconstructions start using ID tracks and their extrapolation to the MS to estimate their momentum at the MS entrance, and three different types of muons are reconstructed in this pass:

- **Segment-tagged (ST) muons** have a full ID track but do not have enough momentum to go beyond the first MS layer. A muon segment is required but is only used for PID, and no combined fit is attempted;

- **Calorimeter-tagged (CT) muons** have a full ID track and a corresponding deposit in the ECAL compatible with a deposit from a MIP. No information from the MS is used;

- **Inside-out combined muons** have a full ID track and muon segments, with their kinematics reconstructed using a full global χ^2 fit. These muons along with CT muons are useful in recovering muon reconstruction efficiency in the $|\eta| < 0.1$ region, where half the sectors are not instrumented and hence an *outside-in* approach fails.

Figure 2.14 illustrates the different types of muons reconstructed by ATLAS. SiF muons are omitted from the diagram for simplicity.

Figure 2.15 summarizes the steps used in the muon reconstruction for both the primary and the inside-out passes in ATLAS.

LRT Muons

A dedicated muon reconstruction pass, functionally identical to the standard pass, is run using all the MS hits in an event along with LRT tracks (instead of the ID tracks from the standard passes) to gain sensitivity to LLP searches with muon final states. This reconstruction chain runs both *outside-in* and *inside-out*, however, only CB, ST, and CT muons are built using the LRT chain. The muons built from this pass are called LRT muons.

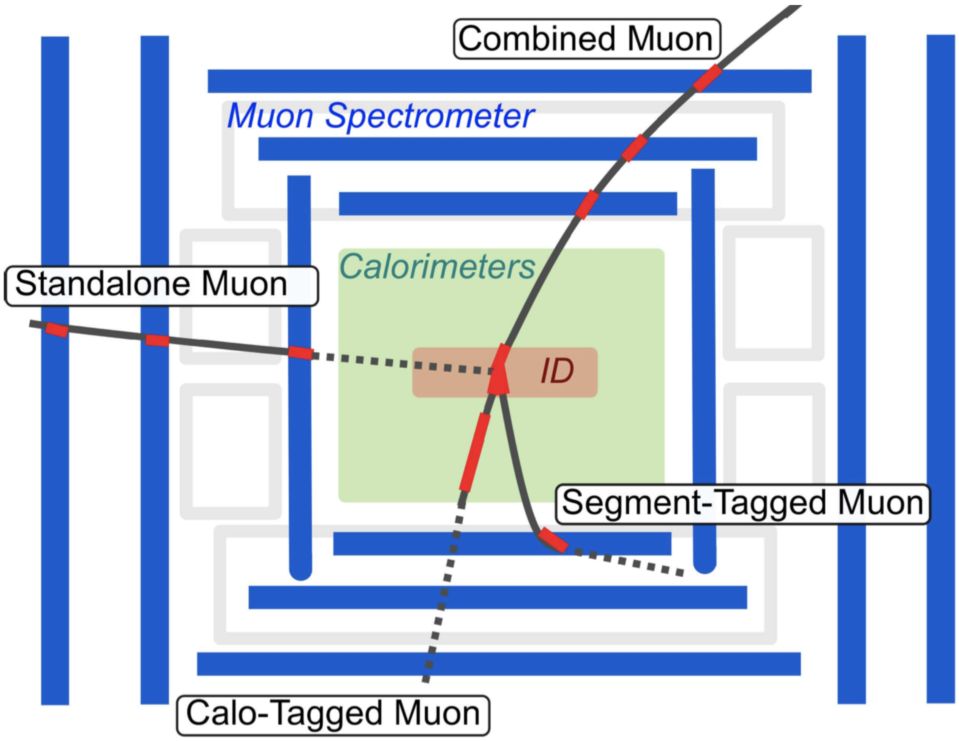

Figure 2.14: The different types of muons reconstructed by the ATLAS experiment. [56]

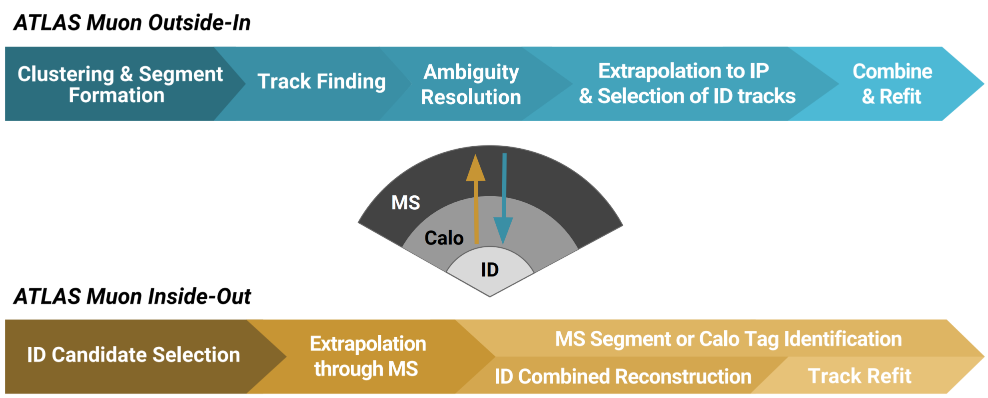

Figure 2.15: Simplified overview of the primary outside–in combined tracking chain and secondary inside–out chain used in ATLAS muon track reconstruction. [3]

2.3.4 Electrons

Electrons in ATLAS are reconstructed using a combination of the ID and the calorimeters [57]. Since electrons can lose a significant amount of their energy because of bremsstrahlung radiation, dedicated calorimetric reconstruction and track reconstruction methods are used to identify and reconstruct them. Figure 2.16 provides a schematic illustration of the elements that enter into the reconstruction of an electron.

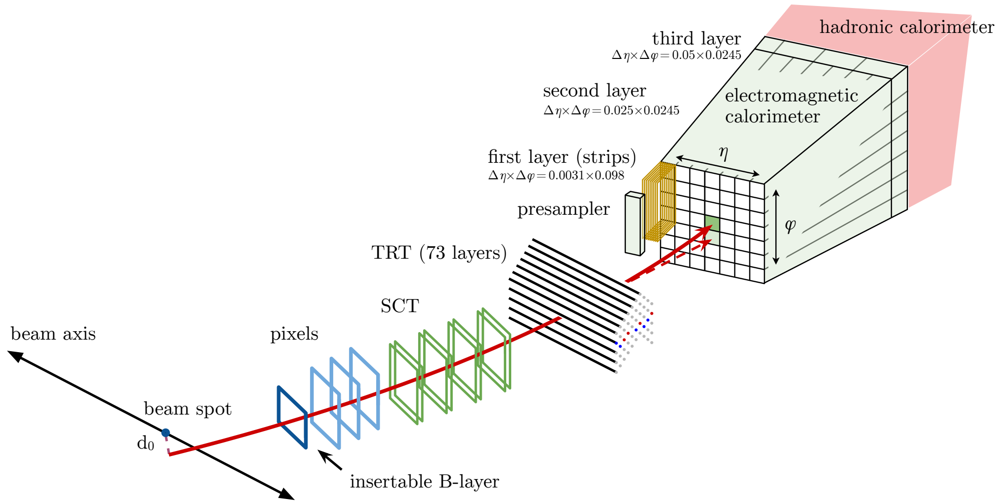

Figure 2.16: A schematic illustration of the path of an electron through the detector. The red trajectory shows the hypothetical path of an electron, which first traverses the tracking system (pixel detectors, then silicon-strip detectors and lastly the TRT) and then enters the electromagnetic calorimeter. The dashed red trajectory indicates the path of a photon produced by the interaction of the electron with the material in the tracking system. [57]

The first step in electron reconstruction is seed-cluster formation. The three calorimeter layers and the pre-sampler are split into a grid in $\eta \times \phi$, and the energy deposited in each point (or *tower*) in the grid is measured. Localized energy deposits are used to form electromagnetic-energy clusters. If two clusters share the same tower, the one with the higher energy is retained.

The second step in the reconstruction is track formation. Unlike MIPs, electrons interact strongly with the ID material and radiate photons. These decay products are expected to be highly collinear with the direction of the outgoing electron. A re-fit of the ID tracks is performed to build tracks for electron candidates allowing for energy losses up to 30% due to bremsstrahlung using a Gaussian sum filter (GSF) [58]. The re-fit, designed to better account for energy loss of charged particles in material, is applied to the ID clusters of raw measurements. This procedure aims to recover tracks loosely matched to EM clusters. The GSF method accounts for radiative looses from electrons, and hence has a better d_0 resolution and q/p accuracy for electrons than those from standard track reconstruction. The tracks built from this algorithm are called GSF tracks.

The last step in the reconstruction is the positional matching of the seed-clusters to the GSF tracks. Multiple ambiguity resolution steps are run to resolve tracks matched to the same clusters, as well as to differentiate genuine electron candidates from mis-identified photons. The energy of the electron is computed from the extended seed-cluster, while the directions are taken from the corresponding GSF track parameters.

LRT Electrons

Similar to muons, a second reconstruction pass for non-pointing electrons is run using all the calorimeter clusters along with GSF tracks refit from LRT candidates (or LRT GSF tracks) to gain sensitivity to LLP searches with electron final states. The electrons built from this pass are called LRT electrons.

2.3.5 Jets

Jets are highly collimated showers of particles created by the rapid hadronization of quarks and gluons. The hadrons in the jets penetrate through the ECAL and are fully absorbed in the HCAL. Many algorithms exist for the reconstruction of jets, however, this analysis uses a particle flow implementation [51] of the anti-k_t algorithm [59]. This implementation uses information from the clusters in the calorimeter as well as tracks in the ID to build jet objects with high energy and momentum resolution.

Anti-k_t is a sequential recombination jet clustering algorithm, in which nearby[5] particles are iteratively collected and combined to form pseudo-jets. This process iterates over all set of particles available in the event till the *size* of the pseudo-jet crosses a pre-defined distance metric (set to be 0.4 for the jets used in this analysis), at which point the pseudo-jet is deemed to be a jet [59].

Figure 2.17 shows a schematic of the particle flow jet reconstruction algorithm used in

[5]the *closeness* of two particles or a particle and a pseudo-jet in this algorithm is determined by their transverse momenta, rapidities and azimuthal angles.

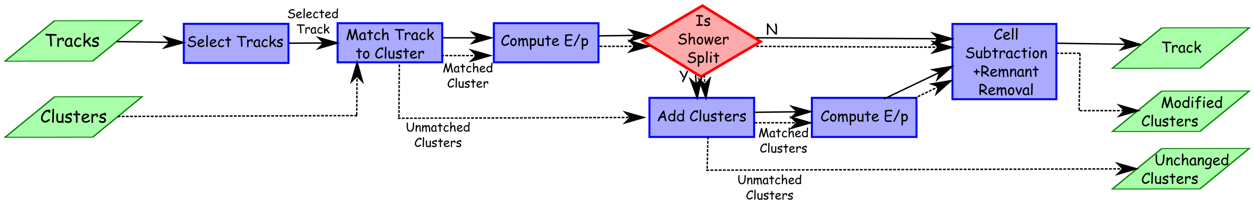

Figure 2.17: A flow chart of how the particle flow algorithm proceeds, starting with track selection and continuing until the energy associated with the selected tracks has been removed from the calorimeter. [51]

ATLAS. Since tracks have better momentum resolution at lower p_{T}, and the calo-clusters have better energy resolution at high transverse energy deposits, the complementary information provided by the two sets of objects results in well-measured hadronic objects, where the ID is used for a measurement of the charged particle kinematics and the calorimeters provide information for the neutral particles. The particle flow algorithm provides a list of selected tracks and a list of topological-clusters (clusters of energy deposits in the calorimeters) containing initially unmodified clusters, as well as a new set of clusters resulting from the energy subtraction procedure described in fig. 2.17. In a case where a cluster and an ID track are matched to one another, the track measurement is used in place of the cluster, and its energy contribution is subtracted from the cluster to yield a modified cluster. Jets reconstructed from this full set of clusters and selected tracks are superior in terms of resolution and pileup robustness.

Chapter 3

ANALYSIS OVERVIEW

The LHC provides a unique opportunity to search for weakly coupled $\mathcal{O}$(GeV) Heavy Neutral Leptons with the ATLAS detector given the large dataset of W-bosons. In 2022, the ATLAS experiment performed a search for HNLs, in the same final state targeted in this analysis, with the full Run-2 dataset [60] (referred to as the 2022 analysis throughout this dissertation). However, the ATLAS object reconstruction algorithms have gone through significant improvement [3] since then, opening up the opportunity to reanalyze the dataset with enhanced techniques and a better overall performance. Section 3.1 highlights the improvements in the reconstruction chain that the analysis benefits from and the corresponding new parameter space that these improvements unlock. Section 3.2 lists out the dataset and the Monte Carlo samples used in this search. Finally, section 3.3 discusses the requirements imposed on the various objects used in this analysis.

3.1 Analysis Goals

This analysis searches for Heavy Neutral Leptons using 140.1 fb^{-1} of pp collision data recorded at $\sqrt{s}$ = 13 TeV in the $W \to \ell_\alpha \mathcal{N}(\to \ell_\beta \ell_\gamma \nu)$ final state. The charged lepton coming from the W decay (ℓ_α) is called the *prompt* lepton since it appears to come from the primary pp collision. Since the HNL is long-lived, the $\ell_\beta \ell_\gamma$ system manifests as a two-pronged Secondary Vertex (SV) or Displaced Vertex (DV)[1], and ℓ_β and ℓ_γ are called *displaced* leptons. The ATLAS ID has the highest granularity (in the $\eta - \phi$ directions) out of all the sub-detector systems, making it the most adept at reconstructing the position and the kinematics of this DV. Hence, the analysis targets a reconstructed two-track DV in the ID volume. The reconstruction selection criteria, or cuts, on the ID tracks (specifically the $|d_0| \leq 300$ mm cut and the minimum number of hits) automatically impose a requirement that the DV radial position (r_{DV}) has to be within the first layer of the strips silicon detectors (see fig. 2.6).

[1]secondary, because the pp collision and the W production and decay gives the primary vertex, and displaced, because the the HNL has a proper lifetime long enough to resolve the PV and the SV.

Furthermore, the SV is required to be displaced ($r_{\rm DV} > 4$ mm), adding a minimum requirement on the radial distance between PV and the DV in the radial direction. Finally, ℓ_α, ℓ_β, ℓ_γ are required to be either μ or e, since τ-leptons are unstable and decay rapidly. The electric charge of ℓ_α follows that of the W boson, while ℓ_β and ℓ_γ are oppositely charged. The goal of this analysis can hence be summarized as:

A search for Heavy Neutral Leptons in decays of W bosons in events with a prompt muon or electron, and a two-track displaced vertex in the Inner Detector volume of ATLAS with outgoing oppositely charged muons or electrons.

3.1.1 Reconstruction improvements

The final state probed in this analysis is non-standard, since the primary ATLAS reconstruction chain is designed around reconstructing signatures from prompt particle decays, as described in section 2.3. Dedicated reconstruction techniques, such as LRT, add significant sensitivity to this final state. The 2022 analysis used an older, computationally expensive version of LRT which was only run on about 10% of the events retained by the ATLAS triggers based on a pre-defined set of filters. Furthermore, the LRT pass used for the 2022 analysis reconstructed a significant amount of "fake" tracks, which inflated the amount of background faced. Lastly, LRT was only run for HNL signal Monte Carlo samples and not for other SM samples, which inhibited the usage of Monte Carlo-driven background estimation techniques.

The analysis presented in this dissertation overcomes all of these challenges thanks to the vast improvement in reconstruction techniques, specifically in LRT. The new optimized LRT deployed for reconstruction in Run-3 [4] was retroactively applied to data collected in Run-2 and all Monte Carlo simulations prepared for Run-2 conditions. The computational costs are significantly lowered, removing the need for filtering of events, making it possible to look for LLP events in all ATLAS data. The optimization significantly reduces the rate of fake tracks (and hence non-LLP SVs) reconstructed by the LRT chain while maintaining high signal efficiency as illustrated in fig. 3.1, which compares the number of reconstructed DVs (both, genuine LLP and fake, i.e. non-LLP) between the legacy (implementation used in the 2022 analysis) and updated LRT reconstruction as a function of $r_{\mathrm{DV-PV}}$ (or L_{xy}) using as a benchmark a simulated sample of $H \to aa \to b\bar{b}b\bar{b}$, where a is an LLP.

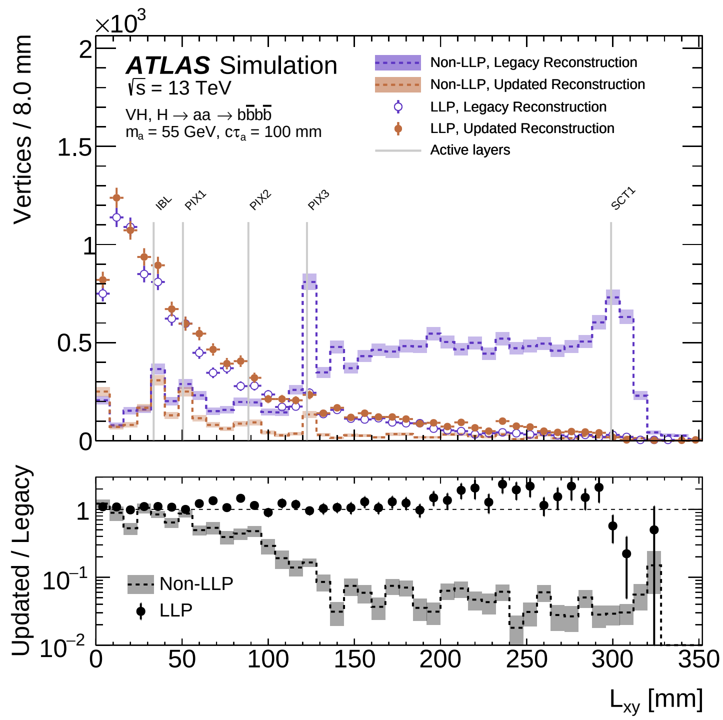

Figure 3.1: A comparison of the radial distributions of reconstructed secondary vertices in a simulated LLP sample using the legacy and updated track reconstruction configurations. [4]

3.1.2 Target Phase Space

The vast improvements in LRT warrant a new search for displaced HNLs in leptonic final states using the full Run-2 dataset. Figure 3.2 summarizes the current limits imposed on the square of the coupling of the HNL to a muon neutrino ($|U_{\mu\mathcal{N}}|^2$) in a single-flavor mixing scenario. The 2022 ATLAS analysis, shown in grey, excluded HNLs in the 1-10 GeV mass range with the corresponding lower limits on $|U_{\mu\mathcal{N}}|^2$ lying in the 10^{-4} and $2{\cdot}10^{-7}$ order. The analysis described in this thesis aims to expand the 2022 exclusion contour in all directions. Specifically, expansion to lower couplings in the 1-3 GeV m_{HNL} range probes HNL parameter spaces untouched by any other experiment. Furthermore, expanding to exclude $m_{\text{HNL}} >10$ GeV with couplings smaller than 10^{-6} allows the exploration of completely new parameters. The former goal is achievable by more advanced background estimation methods, since the $\mathcal{O}$(GeV) m_{HNL} regions are generally dominated by background from hadronic decays. The latter goal is achievable by exploiting the improved LRT and DV reconstruction, since that phase space is limited by background with DVs from uncorrelated track crossings.

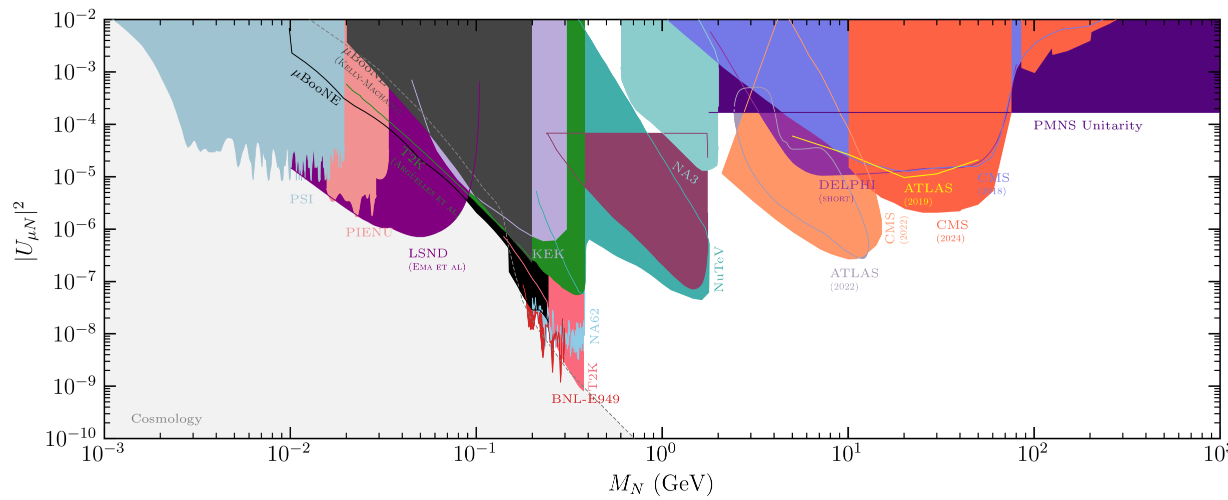

Figure 3.2: Summary of exclusion limits imposed by various experiments on the square of the coupling of Heavy Neutral Leptons to the muon neutrino as a function of the HNL mass. The limits imposed by the ATLAS 2022 analysis is shown in grey as ATLAS (2022). [61]

3.2 Dataset and Monte Carlo Simulations

3.2.1 ATLAS dataset

The full Run-2 ATLAS dataset used for this analysis corresponds to an integrated luminosity of 140.1 fb^{-1} of pp collisions at $\sqrt{s}$ =13 TeV. Only the subset of events (satisfying the Good Run Lists) are considered which have passed detector condition and data quality thresholds to ensure the data are suitable for physics analysis. The amount of data collected is roughly correlated to the pileup during a set period, since the collision rate is larger when a higher number of protons are squeezed together in a bunch.

The ATLAS data goes through a reconstruction chain as described in section 2.3, and is stored in a data format called Analysis Object Data (AOD). These files, due to the presence of low-level event information, are large in size and are designed to be used across ATLAS analyses. A further skimming is applied on these files based on specific analysis needs. Given the non-standard nature of physics objects and reconstruction methods used in this analysis, a specialized data format, called DAOD_LLP1, is used. This Derived AOD (DAOD) format adds information for a large breadth of analyses searching for LLPs which wouldn't be available in the standard DAOD formats. The secondary vertexing algorithm used in this analyses requires low-level track information which are available in the AODs but not in the DAODs, and hence is run at the AOD→DAOD_LLP1 step.

3.2.2 Simulated Samples

The signal and background processes affecting this analysis are simulated using Monte Carlo (MC) event generation. These simulations start with the event generation of the underlying hard-scatter event along with the parton shower process. The simulation is followed by the modeling of the interaction of the particles with the active detector material done using the Geant4 [62] simulation toolkit. In the next step, additional *pp* interactions are overlaid along with the hard-scatter according to the data taking conditions of a corresponding year, and the whole event goes through a digitization step where the detector readout response is simulated to represent the signature left by real particles in ATLAS. Three MC campaigns are simulated - mc20a, d, and e, corresponding to the pileup and detector conditions in the years 2015-16, 2017, and 2018, respectively. The simulated pileup profile is reweighted to match the total integrated luminosity of the data taking period. The digitized signature is reconstructed to build physics objects in the AOD format. Finally, DAOD_LLP1 are created for both data and MC samples to ensure the same treatment is applied to both.

Signal MC Samples

Signal samples were generated using MadGraph5_aMC@NLO [63] with HeavyN UFO model libraries [64] that provide a minimal SM extension with HNLs at next-to-leading order QCD accuracy and uses the NNPDF3.0nlo [65] set of PDFs. The events were interfaced to Pythia 8.230 [66] to model the parton shower, hadronization, and underlying-event, with parameters set according to the A14 tune [67] and using the NNPDF2.3nlo [68] PDFs.

The W boson from the primary pp interaction is required to decay into a muon or an electron, and an HNL admixture ($W \rightarrow \ell\mathcal{N}$). The HNL is given a long lifetime and set to decay leptonically ($\mathcal{N} \rightarrow \ell\ell\nu$, $\ell = \mu$ or e, ν is the SM neutrino). The decay is modeled via a V-A weak matrix element, and can be mediated either via a virtual W^* or Z^*. Both diagrams contribute to the decay if the charged leptons in the HNL decay have the same flavor. Samples are created with mean decay lengths $c\tau_{\text{HNL}} = 1$, 10, 100 mm and masses $m_{\mathcal{N}} = 1$, 2, 3, 4, 5, 7.5, 10, 12.5, 15, 17.5, 20 GeV. For masses up to 5 GeV, $c\tau_{\text{HNL}} = 1000$ mm is also considered. Both lepton-number conserving and violating decays are simulated.

The various models considered in this analysis are summarized in table 1.2. The simulated samples are scaled to different cross-sections based on the final state and the model being considered. Six decay modes are considered in this analysis determined by the flavor of the charged leptons. The channels are represented as $\ell_\alpha - \ell_\beta\ell_\gamma$, where α is the flavor of the prompt lepton, and β and γ are the flavors of the charged leptons from the HNL decay. Table 3.1 summarizes the six final states considered in this analysis and the decay modes that lead to that final state. ℓ_β and ℓ_γ cannot be experimentally resolved and neutrinos flavors cannot be observed using the ATLAS detector. Hence, the $\ell_\beta\ell_\gamma\nu$ and $\ell_\gamma\ell_\beta\nu$ contributions are combined.

Background MC Samples

MC simulations are used to model the kinematics of the (correlated) background from SM processes contaminating the phase space explored in this analysis. Specifically, simulations of the top-quark pair production ($t\bar{t}$) and V+jets ($V = W, Z$) processes are considered.

Final state/ Channel	Decay mode from model with Single flavor mixing HNL	$\mu-$ and $e-$ mixing HNL
$\mu - \mu\mu$	$W \to \mu\mathcal{N}\,(\to \mu\mu\nu_\mu)$	$W \to \mu\mathcal{N}\,(\to \mu\mu\nu_\mu + \mu\mu\nu_e)$
$\mu - \mu e$	$W \to \mu\mathcal{N}\,(\to \mu e\nu_e)$	$W \to \mu\mathcal{N}\,(\to \mu e\nu_e + e\mu\nu_\mu)$
$\mu - ee$	$W \to \mu\mathcal{N}\,(\to ee\nu_\mu)$	$W \to \mu\mathcal{N}\,(\to ee\nu_\mu + ee\nu_e)$
$e - ee$	$W \to e\mathcal{N}\,(\to ee\nu_e)$	$W \to e\mathcal{N}\,(\to ee\nu_e + ee\nu_\mu)$
$e - e\mu$	$W \to e\mathcal{N}\,(\to e\mu\nu_\mu)$	$W \to e\mathcal{N}\,(\to e\mu\nu_\mu + \mu e\nu_e)$
$e - \mu\mu$	$W \to e\mathcal{N}\,(\to \mu\mu\nu_e)$	$W \to e\mathcal{N}\,(\to \mu\mu\nu_e + \mu\mu\nu_\mu)$

Table 3.1: The six final states considered in this analysis and the decay modes that contribute to them for the different mixing models considered.

The production of $t\bar{t}$ events was modelled using the POWHEG BOX v2 [69, 70, 71, 72] generator at next-to-leading order with the NNPDF3.0NLO [65] PDF set and the h_{damp} parameter[2] set to $1.5\,m_{\mathrm{top}}$ [73]. The events were interfaced to PYTHIA 8.230 [66] to model the parton shower, hadronisation, and underlying event, with parameters set according to the A14 tune [67] and using the NNPDF2.3LO set of PDFs [68]. The decays of bottom and charm hadrons were performed by EVTGEN 1.6.0 [74]. The non-all-hadronic slice of the sample is used, which includes leptonic final states.

The $W(\to \ell\nu)$+jets and $Z(\to \ell\ell)$+jets samples are generated using SHERPA2.2.11 [75, 76, 77] at next-to-leading order accuracy up to 2 jets, and up to leading order accuracy for up to 5 jets in the Matrix Element. The samples are created in three slices of the flavor of the jets, namely BFilter, CFilterBVeto, and CVetoBVeto, where XFilterYVeto refers to samples where a jet with flavor X is required and a jet with flavor Y is vetoed. Furthermore, for τ final states, the samples are divided by the decays of the τ lepton as either leptonic or (semi-)hadronic. For the predictions used in this analysis, all the slices and filters are combined for the total expected event yield.

[2]The h_{damp} parameter is a resummation damping factor and one of the parameters that controls the matching of POWHEG matrix elements to the parton shower and thus effectively regulates the high-p_{T} radiation against which the $t\bar{t}$ system recoils.

3.2.3 Signal Kinematics

The position of the HNL decay depends on its proper lifetime and its boost, and the kinematics of its decay products depend on its mass and its boost. Figure 3.3 shows some representative kinematic distributions of the HNL and its decay products at the *truth level*, i.e. the actual kinematics of the simulated particles inside the ATLAS experiment. Such distributions are typically contrasted against kinematics at the *reconstructed level*, i.e. kinematics after digitization, reconstruction, and identification of the same simulated particles.

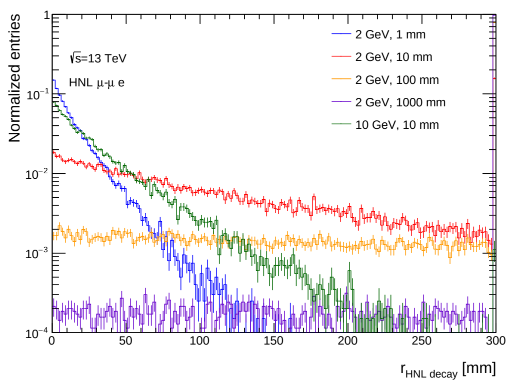

(a) HNL decay radius

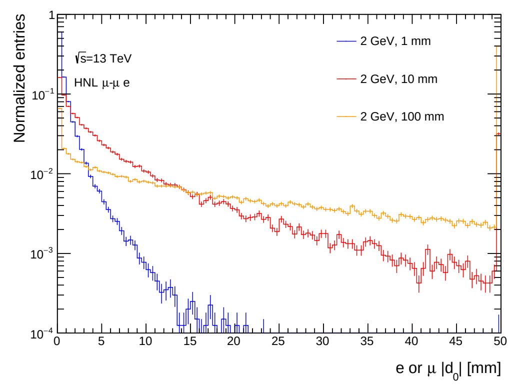

(b) μ or e $|d_0|$

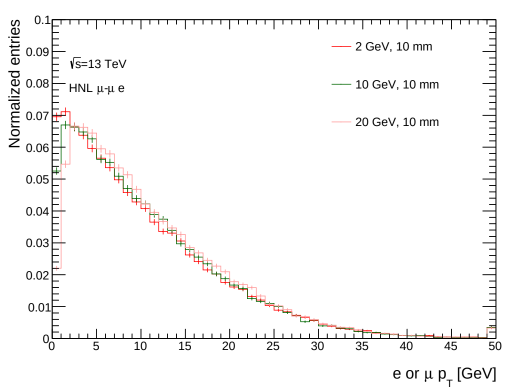

(c) μ or e p_{T}

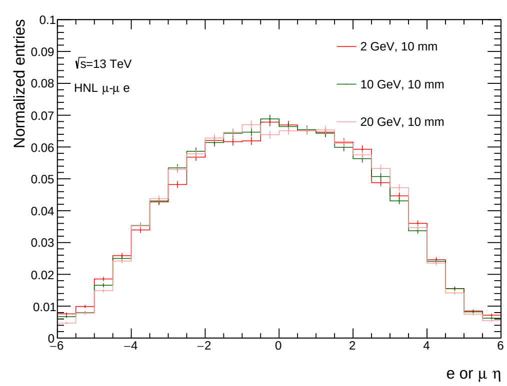

(d) μ or e η

Figure 3.3: Truth level kinematic distributions for $\mu - \mu e$ HNL samples for some representative masses and proper lifetimes indicated in the legend. The total entries for each histogram are normalized to 1 for a fair comparison comparison, and overflow (underflow) is shifted into the first (last) bin.

In the frame of the HNL, the probability density for it to decay after a time t is given by the exponential distribution:

$$p_{\text{decay}}(t) = e^{-t/\tau_{\text{HNL}}}. \tag{3.1}$$

In the lab frame, the radial position of the HNL decay, as illustrated in fig. 3.3a, depends on both its mass and proper lifetime due to a time dilation effect from its Lorentz boost. The track reconstruction efficiency drops after 300 mm due to the lack of enough silicon hits in the ID, causing most of the HNL decays for $c\tau_{\text{HNL}}$= 100 mm and 1000 mm to lie outside acceptance. Figure 3.3b shows the $|d_0|$ distribution for the charged decay products of the HNL which is independent of the mass of the HNL for a fixed proper lifetime. LRT adds sensitivity to higher lifetime HNLs since they have a large proportion of decays with $|d_0| > 5\,\text{mm}$. Figures 3.3c and 3.3d show that the p_{T} and η spectrum of the HNL decay products have a weak dependency on the mass and the lifetime being considered, especially in the $m_{\text{HNL}} << m_W$ regime, where most of the momentum from the W decay is transferred to the prompt lepton.

After studying the truth-level kinematics of the HNL signal, the simulated samples are used to quantify the detector acceptance and reconstruction efficiencies of the HNL samples as shown in fig. 3.4.

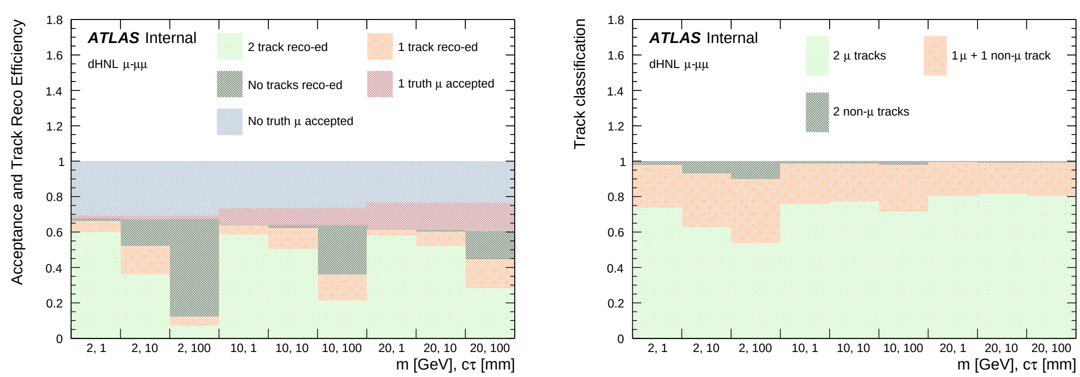

(a) Acceptance Categorization

(b) Reconstruction Categorization

Figure 3.4: Fractional detector acceptance, ID track reconstruction, and muon reconstruction efficiencies for $\mu - \mu\mu$ HNL samples for some representative mass and proper lifetimes indicated on the $x-$axis.

Figure 3.4a classifies HNL di-muon decays according to the detector and kinematic acceptance ($p_{\mathrm{T}} > 3$ GeV, $|\eta| < 2.7$) of a truth muon and if the truth muon leaves hits that are reconstructed as an ID track. The fraction of events where both muons are reconstructed as tracks is within 5-60% with a strong dependence on the HNL proper lifetime and a weak dependence on its mass, since they affect the production radius of the muons and hence their tracking acceptance. Figure 3.4b further categorizes the cases with two reconstructed ID tracks according to their muon extension efficiency i.e. if the ID tracks are matched to reconstructed muons. A stronger (yet still subtle in the masses considered) dependency is observed on the HNL mass than lifetime, since the p_{T} spectrum of the muons are skewed towards higher values for heavier HNLs, as shown in fig. 3.3c. Hence, low mass and high lifetime HNLs suffer from low detector acceptance and low lepton reconstruction efficiencies. This fact, further convoluted with the low cross-section of HNL production in that phase space, makes the exploration of lighter and weakly-coupled HNLs extremely challenging.

3.3 Physics Objects Selection

The final state probed in this analysis is characterized by a prompt lepton and a displaced vertex in the ID volume with two outgoing leptons. Furthermore, some requirements on the quark flavor of jets in the event are imposed in the definition of analysis regions.

3.3.1 Prompt Lepton and Trigger

A prompt lepton is defined as a muon or an electron coming from the primary pp interaction. The prompt lepton is used as the primary triggering object to tag the W-decay, and hence, the requirements on the selection of the prompt lepton are driven by the selections imposed by the trigger chain. The identification Working Point (WP) adds self-consistency and goodness of fit requirements on the lepton object to select genuine leptons against incorrect combinations of hits or hadron decays-in-flight. Similarly, isolation separates leptons without hadronic or radiative activity around them, as is expected from $W \to \ell\nu$, from those that are within jets. Finally, track-to-vertex association requires that the lepton points back to the primary pp collision. Table 3.2 summarizes the selections imposed on the prompt lepton candidate.

Cut Name	Cut Description
Lepton Flavor	muon or electron
Trigger Matched	✓
Transverse Momentum	$p_{\mathrm{T}} > 27$ GeV
Identification Working Point	`Medium` [55, 57]
Isolation Working Point	`PflowLoose_VarRad` [55] (`Loose_VarRad` [78]) for $\mu\,(e)$
Track-to-Vertex-Association	$\lvert d_0/\sigma(d_0)\rvert < 3\,(5)$ for $\mu\,(e)$, $\lvert z_0 \sin\theta\rvert < 0.5$ mm [55, 78]

Table 3.2: Selections used for the prompt lepton candidate.

If more than one prompt lepton candidate are found in an event, the candidate with the highest transverse momentum is chosen as the prompt lepton.

Period	Muon Trigger	Electron Trigger
2015	HLT_mu20_iloose_L1MU15	HLT_e24_lhmedium_L1EM20VH
2016, 2017, 2018	HLT_mu26_ivarmedium	HLT_e26_lhtight_nod0_ivarloose

Table 3.3: Lowest unprescaled single lepton triggers used in the analysis.

The lowest unprescaled single lepton triggers are used in this analysis, with the exact trigger scheme for each data taking year listed in table 3.3. Figure 3.5 shows the fraction of events where a single lepton fires the trigger as a function of the prompt lepton p_{T} for the single muon and single electron triggers for a selection of $\mu - \mu\mu$ and $e - ee$ signal samples. The low p_{T} tail on the plots come from events in which the lepton causing the trigger originated from the displaced leptons. This tail would disappear once the selections on the prompt lepton are applied. The selections in fig. 3.5 are limited to the presence of a PV, a prompt lepton (without any selections), and the presence of a $\mu\mu$ or ee DV in the fiducial volume.

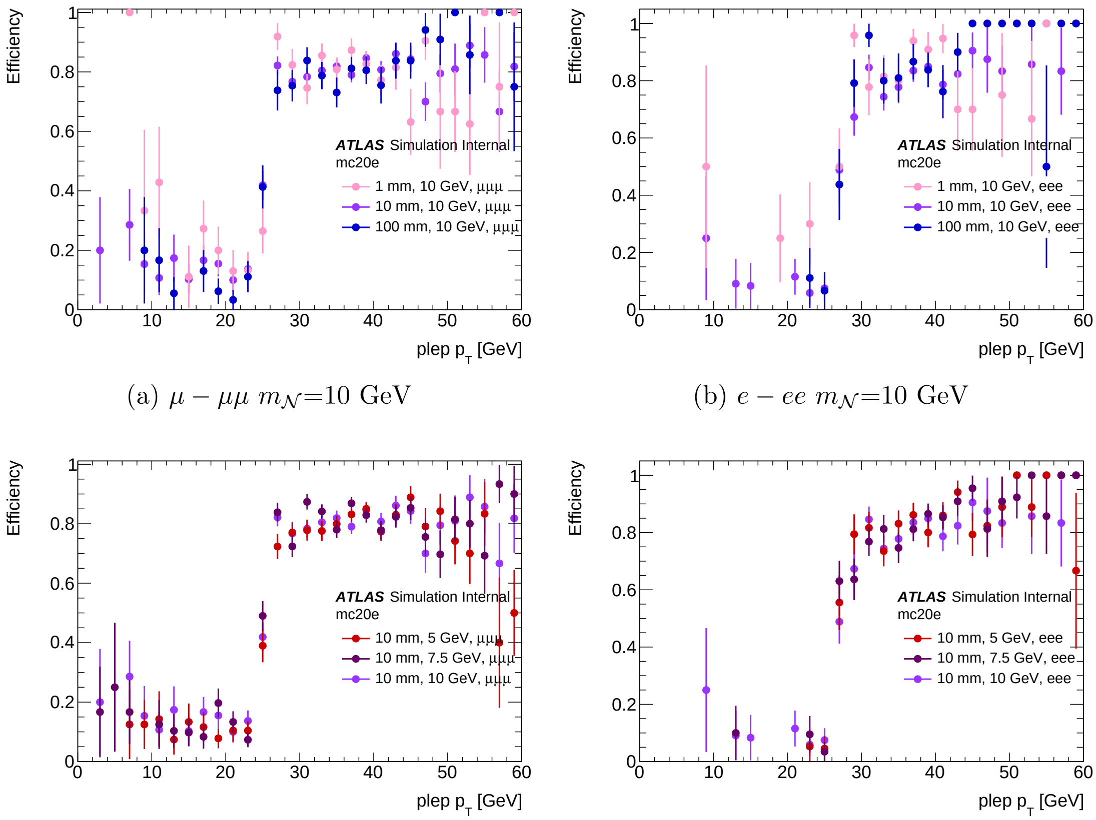

(a) $\mu - \mu\mu$ $m_{\mathcal{N}}$=10 GeV

(b) $e - ee$ $m_{\mathcal{N}}$=10 GeV

(c) $\mu - \mu\mu$ $c\tau_{\mathrm{HNL}}$=10 mm

(d) $e - ee$ $c\tau_{\mathrm{HNL}}$=10 mm

Figure 3.5: Fraction of events where any single-lepton trigger is fired by any lepton as a function of the prompt lepton p_{T} for a selection of HNL signal mass and lifetimes for the $\mu - \mu\mu$ and $e - ee$ channels.

3.3.2 Displaced Muons

Muons in the displaced vertex are referred to as displaced muons. Since these muons are produced from a decay away from the IP, they are not always expected to point back to it. Hence, to extend the analysis coverage to such non-pointing muons, a combination of standard and LRT muons are used.

As described in section 2.3, the standard (i.e. both primary and inside-out) muon reconstruction and LRT muon reconstruction run seperately using all the hits in the MS in an event as inputs along with the standard ID tracks for the former chain and the LRT tracks for the latter chain. This workflow, due to the use of all MS hits in both the chains, causes a large fraction of LRT muons to share MS hits (or rather a full MS track) with a standard muon. The LRT and standard muon are hence said to have an *overlap* with each other. This is an unphysical situation since it's not possible for two genuine muons to leave identical MS hits. However, the muon reconstruction algorithm is very liberal in the quality of muons it reconstructs, and hence a bad quality LRT muon could overlap with a good quality standard muon, and vice versa. An overlap resolution strategy is designed to retain or reject a muon object from a pair (out of which one is reconstructed as a standard muon and the other as a LRT muon) which share an MS track between them. The strategy is optimized for maximal signal retention post overlap removal, and uses fundamental muon quality requirements without adding any biases from the kinematics of the signal. In case both or neither muon passes a check, the resolution moves on to the next check. The overlap resolution is designed to, in case of a MS track overlap, retain the muon which passes these requirements, followed sequentially:

- passes the `Loose` [55][3] identification working point,
- is a Combined muon,
- has a lower absolute η difference between the measurement from ID track and the MS track extrapolated to the perigee,
- is reconstructed in the standard pass (fallback option).

Figure 3.6 shows the fraction of standard and LRT muons which have an overlap at different stages of the overlap removal process.

After the standard and LRT muon collections are combined and the overlaps are resolved, they are treated similarly for downstream analysis. Only Combined displaced muons are used in this analysis with the kinematic requirements of $p_{\mathrm{T}} > 3$ GeV and $|\eta| < 2.5$, both of which are imposed by muon reconstruction and detector constraints. A `Medium` identification WP is used to identify genuine muons against pion decays. However, the application of the WP

[3]The `Loose` criteria used here is modified to not apply any requirements on the ID track.

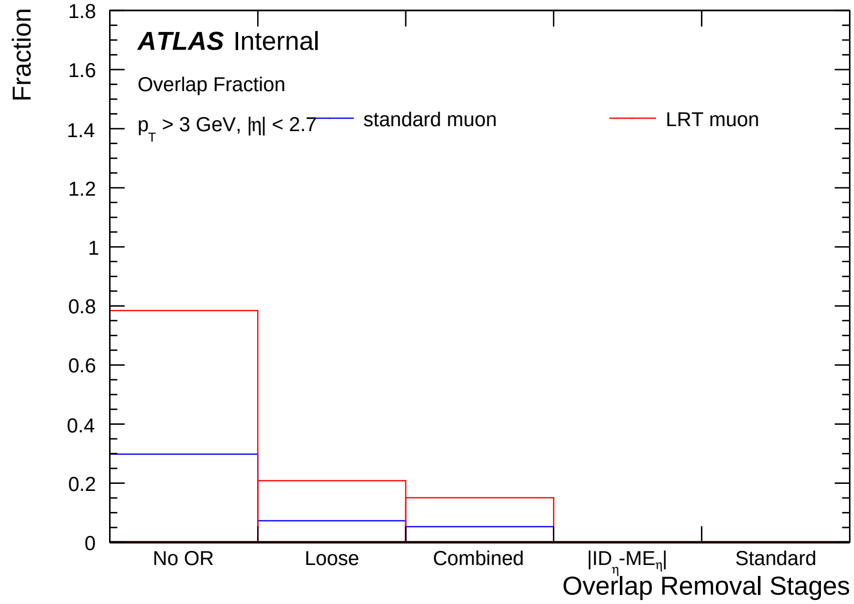

Figure 3.6: Fraction of standard and LRT muons that have MS track overlaps at different stages of the overlap removal process.

is modified for LRT muons. LRT reconstruction has, by design, stringent reconstruction criteria as shown in table 2.3. Hence, the `Medium` ID WP for LRT muons is slightly modified to not apply the above requirements on the ID tracks while retaining the other cuts. Table 3.4 lists the cuts imposed on the identification of displaced muons.

Cut Name	Cut Description
Transverse Momentum	$p_{\mathrm{T}} > 3$ GeV
Pseudo-rapidity	$\lvert\eta\rvert < 2.5$
Muon Type	Combined
Identification Working Point	`Medium` (w/o ID cuts for LRT muons)

Table 3.4: Selection criteria used for displaced muons.

3.3.3 Displaced Electrons

Electrons in the displaced vertex are referred to as displaced electrons. Similar to muons, a combination of standard and LRT electrons are used to identify displaced electrons. The standard and LRT electron reconstruction use the same calorimeter information for each event, and hence, a standard electron and an LRT electron can point to to the same calo-cluster. These *overlaps* are resolved by using the trigger-level identification WPs as defined in [57]. Out of an overlapping standard-LRT electron pair, the electron which passes a stricter[4] identification WP requirement is retained. In the case where both the standard and LRT electron pass the same identification WP, the electron reconstructed from the standard pass is kept. The identification WPs used here are further loosened with respect to the trigger-level working points by relaxing the hard requirements on the number of silicon hits in the pixel sub-detectors.

A p_{T} threshold of 4.5 GeV is imposed on displaced electrons driven by reconstruction requirements. No cuts beyond those imposed by detector limitations are applied on electron η. The `VeryLooseNoPix` identification WP is used for the displaced electrons, which is the same as the `VeryLooseNoPix` trigger-level identification WP as defined in [57] (which already uses a relaxed version of the WPs used for offline electron identification), but with the requirement on the number of silicon pixel hits removed to reduce the bias towards electrons emerging from the IP and add efficiency to electrons from LLP decays. The same identification WP is applied to both standard and LRT electrons. Since the WP imposes very relaxed requirements on the electron, an extra requirement on the compatibility of the p_{T} of the electron object and its corresponding GSF track is imposed to reject mismatches that would arise from photon radiation. Such a cut ensures that the final state identification, which is based on the electron object, and the DV kinematics reconstruction, which is based on the track object, are compatible to each other. Table 3.5 lists the cuts imposed on the identification of displaced electrons.

Cut Name	Cut Description
Transverse Momentum	$p_{\mathrm{T}} > 4.5$ GeV
Identification Working Point	`VeryLooseNoPix`
Momentum Balance	$\lvert p_{\mathrm{T}}^{\mathrm{electron}} - p_{\mathrm{T}}^{\mathrm{GSF\ track}} \rvert / p_{\mathrm{T}}^{\mathrm{electron}} < 0.5$

Table 3.5: Selection criteria used for displaced electrons.

[4]The `Tight` identification WP has the strictest requirements, followed by `Medium`, `Loose`, and `VeryLoose`.

3.3.4 b-tagged Jets

At tree level, the HNL signal process in the fully leptonic final states has no hadronic activity. The leading source of jets in signal events is Initial State Radiation from the incoming partons which then showers to give jets. The chance of hence observing a jet (which have a reconstruction p_{T} threshold of 20 GeV) is low, and a jet from a b-quark is even lower. This fact is used to define analysis regions rich in background and highly depleted in the signal process.

This analysis uses anti-k_t $\Delta R = 0.4$ PFlow jets reconstructed with $p_{\mathrm{T}} > 20$ GeV and $|\eta| < 4.9$ limited by the calorimeter coverage. Machine learning algorithms are trained to learn about the substructure of jets to classify jets initiated by heavy (b or c) quarks against those initiated by light quarks or gluons. In this analysis, the DL1d tagger [79] is used to identify the flavor of jets. DL1d is based on a deep neural network. It takes information about the secondary vertex reconstructed inside jets and other low level information about jet substructure and returns a classification score. A cut on the DL1d tagger score of a jet which accepts on average 85% of jets with b-hadrons, or the 85% b-jet efficiency WP, is applied to identify jets as either coming from b-quarks, or lighter quark flavors. A jet passing the 85% WP cut of the DL1d tagger score is called a b-tagged jet or a b-jet.

3.3.5 Overlap Removal

This analysis makes use of reconstructed electron, muon, and jet physics objects. Since the reconstruction chains of these objects run in parallel, it is possible that such objects share ID tracks or calo-clusters. An overlap removal (where the overlap criteria based on the type of objects in context) step is run after the baseline selection of the physics objects but before using them for analysis. An overlap removal strategy which favors leptons over jets is used in this analysis. This choice was shown to have a negligible impact on the signal significance in the measured phase space as compared to the standard ATLAS overlap removal strategy. However, a significant increase in the event yield was seen in a non-overlapping analysis region which is used to constrain correlated background contribution. The non-overlapping region is used to collect events with leptons inside jets the statistics of which are increased by using this overlap removal strategy, which led to a better data-driven estimate of the background yield. Table 3.6 lists the rules used for resolving overlaps between physics objects according to the criteria of overlap between them, run sequentially from top to bottom.

Reject	Against	Overlap Criteria
electron^1	electron^2	shared ID track, $p_{\mathrm{T}}^1 < p_{\mathrm{T}}^2$
muon	electron	is CT-muon and shared ID track
electron	muon	shared ID track
jet	electron	$\Delta R < 0.2$
jet	muon	#tracks in jet <3 and (ghost-associated[5]or $\Delta R < 0.2$)

Table 3.6: The lepton-favored overlap removal scheme and overlap criteria.

3.3.6 Secondary Vertex Reconstruction and Selection

The track selection, DV reconstruction, and DV selection is facilitated by a package called Vertex Secondary Inclusive (VSI) [80] at the core of which is the vertex reconstruction algorithm, VkalVrt [81].

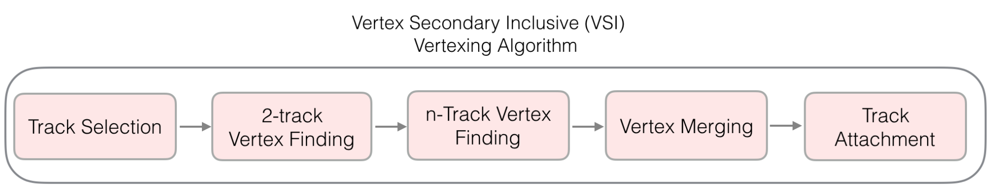

Figure 3.7: The main steps used for secondary vertex reconstruction in the Vertex Secondary Inclusive algorithm. [80]

Figure 3.7 summarizes the steps followed by VSI for DV reconstruction. The first step of the algorithm is the selection of tracks used to seed the reconstruction. A large number of tracks arise from pileup in the current LHC conditions which are not of physics interest for searches such as ours. Hence, a strict selection of tracks is important to reduce the

[5]Particles are assigned to jets following the ghost-association procedure which consists of assigning them to jets by adding their tracks with infinitesimal p_{T} to the jet clustering process. A particle is hence ghost-associated to a jet if it is retained after the jet clustering.

combinatorics and hence chances of reconstructing a DV which comes from a fake crossing rather than a genuine decay. These pre-selected tracks are used to create pairs in a two-track seed finding which checks their compatibility as DV track candidates, followed by a χ^2 minimizing fit. Since a track can be attached to multiple seeded pairs, such cases are resolved in the next step by checking for incompatibilities of a track in an n-track vertex, based on a goodness of fit score. Multiple secondary vertices corresponding to a single LLP decay can be reconstructed by this algorithm. To combine these candidates, DVs with low track multiplicity are attempted to be refit with DVs with higher track multiplicities requiring the refitted DV position and the original DV position to be compatible. In the final step, tracks that are not part of any DVs are attempted to be attached to the reconstructed DVs if they increase the goodness of fit score.

This analysis uses the same steps and configurations as outlined in [80], apart from a few but significant notable differences which result in improved signal acceptance and significant reduction in uncorrelated background:

- Since the final state probed in this analysis is fully leptonic, the track pool used for seeding is modified to use

  1. ID tracks matched to CB or ST muons,
  2. GSF tracks matched to electrons, and
  3. other non-leptonic ID tracks

  in each event. GSF tracks have better d_0 and z_0 resolution than standard ID tracks for electrons and hence result in better reconstruction efficiency for DVs with electrons. Other non-leptonic ID tracks are added to all events to expand sensitivity to lepton+track like final states which would result from semi-leptonic decays of the HNL which are planned to be studied in the near future.

- At the two-track vertex finding level, the minimum $|d_0|$ requirement on the pair of tracks is reduced from 2 mm (as is in the standard algorithm) to 1 mm. An error was found in the technical implementation of this cut which led to the accidental acceptance of a significant number of DVs failing this requirement in the 2022 analysis. This error was fixed for this version of the analysis, leading to a significant reduction in correlated background and a near elimination of uncorrelated background.

- The track attachment is modified to only use tracks from the seeding pool. This reduces the unwanted promotion of genuine two-track signal vertices to three-track vertices with one accidental track crossing, say from pileup. Such a change increases signal acceptance since, in the analysis, n-track vertices are vetoed, where $n > 2$.

- As a last stage cleaning, all DVs without any tracks matched to leptons are removed.

The reconstruction efficiency of DVs is defined as the rate of having a reconstructed vertex matched to an individual HNL decay. To study the tracking and vertexing performance separately from each other, the total efficiency is factorized into independent terms. The efficiencies are defined as follows:

- **Acceptance,** $\mathcal{A}$: fraction of all HNL decays with two charged particles with $p_{\mathrm{T}} > 1$ GeV, a true LLP decay z position of $|z| < 2720$ mm (inner surface of the last endcap layer), and a transverse distance from the origin of $r_{\mathrm{DV}} < 563$ mm (inner surface of the TRT). The first requirement ensures that the vertex has decay products have large enough momentum to be reconstructed by the tracking algorithm, while the latter two ensure that the HNL was produced within the tracking volume of the ID. The definition of acceptance is purely based on the geometry of the ID without taking into account the limits of track reconstruction, which become more relevant in the next term, the seed efficiency.

- **Seed efficiency,** $\varepsilon_{\mathrm{seed}}$: fraction of accepted HNL decays having two reconstructed tracks passing the ID track selection requirements. The seed efficiency determines the performance of the object reconstruction algorithms convoluted with the effect of the track seeding cuts.

- **Core efficiency,** $\varepsilon_{\text{core}}$: fraction of seeded HNL decays having a matched reconstructed vertex. The matching of ID tracks to truth particles is based on a weighted scoring of shared detector hits between the track and truth particle trajectories. For each pair of a reconstructed vertex v and a truth HNL decay vertex l, a truth-matching score s is computed using the magntitude of the track p_{T} as a weight. The score is given by:

$$s(v,l) \equiv \frac{\sum_{i\in \text{ tracks } \in v} \left(p_{\text{T}}^{(i)} \mid \text{ descendent of HNL decay } l\right)}{\sum_{i\in \text{ tracks } \in v} p_{\text{T}}^{(i)}} \tag{3.2}$$

  A reconstructed vertex v is said to match to a truth vertex l if the score $s(v,l) > 0.5$. The core efficiency represents the true performance of the vertexing algorithm itself with the track acceptance and reconstruction effects factored out.

The **total efficiency** ($\varepsilon_{\text{total}}$), is given by

$$\varepsilon_{\text{tot}} = \mathcal{A} \cdot \varepsilon_{\text{seed}} \cdot \varepsilon_{\text{core}}. \tag{3.3}$$

Figure 3.8 shows the four efficiencies as a function of r_{DV}, the radial distance of the HNL decay from its origin for a $m_{\mathcal{N}}$=10 GeV, $c\tau_{\text{HNL}}$=10 mm benchmark HNL signal sample for the three DV flavor combinations $\mu\mu$, $e\mu$, and ee under 2018 pileup and detector conditions. The flavor of the prompt lepton is irrelevant for this study. $\mathcal{A}$ is measured to be uniformly high at ∼80% for all three flavors and falls off at high r_{DV} values as the HNL starts to decay outside the ID acceptance. $\varepsilon_{\text{seed}}$ is highest for $\mu\mu$ followed by $e\mu$ and ee. The p_{T} threshold for electron reconstruction is higher than that for muons (4.5 GeV as compared to 3 GeV), making $\varepsilon_{\text{seed}}$ lower for LLPs decaying to electrons. Furthermore, a failure in reconstructing any of the two charged particles from the HNL decay would lead to a non-seeded vertex. This, along with the stringent requirements imposed by VSI on track quality, results in relatively low $\varepsilon_{\text{seed}}$ across DV flavors and the full r_{DV} range, with a sharp drop to 0 at $r_{\text{DV}} > 300$ mm due to a drop in track reconstruction efficiency. $\varepsilon_{\text{core}}$ is similar for all DV

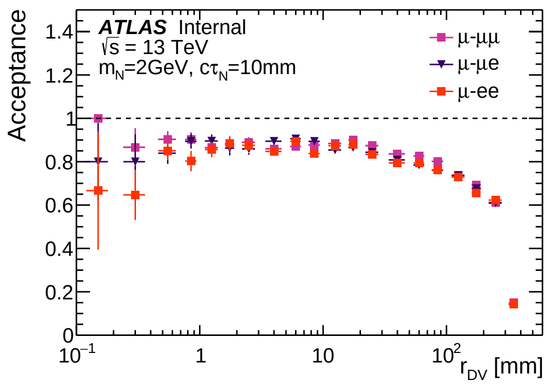

(a) Acceptance

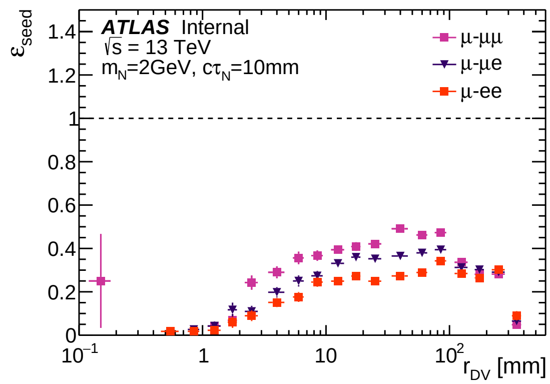

(b) Seed efficiency

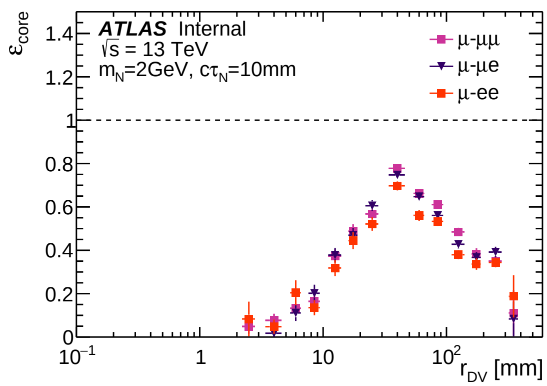

(c) Core efficiency

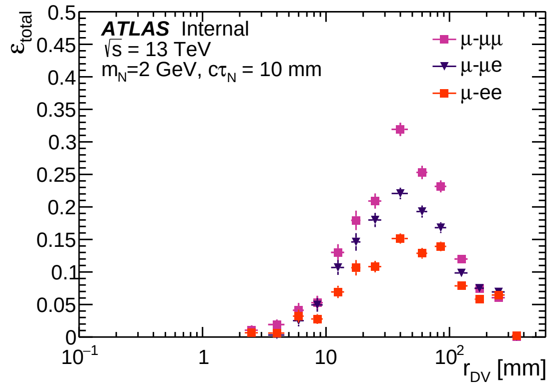

(d) total efficiency

Figure 3.8: The four efficiencies with the custom vertexing configuration as a function of r_{DV} evaluated for the $\mu - \mu\mu$, $\mu - \mu e$, and $\mu - ee$ channels with $m_{\mathcal{N}} = 2$ GeV, $c\tau_{\mathrm{HNL}} = 10$ mm.

flavors. The losses at low r_{DV} are driven by the cut on minimum track $|d_0|$. Furthermore, $\varepsilon_{\mathrm{core}}$ starts dropping at high r_{DV} since, for a given mean proper lifetime, the average HNL boost increases with decay radius, resulting in more collimated decay products and a vertex topology which is more challenging to reconstruct. $\varepsilon_{\mathrm{total}}$ is hence seen to cap at 15-35% for $m_{\mathcal{N}}$=10 GeV, $c\tau_{\mathrm{HNL}}$=10 mm HNL decays driven by the DV flavor.

Apart from the effects mentioned above, the DV reconstruction performance also depends on the pileup since a higher incidence of pileup tracks makes it difficult to separate genuine decays from decays from accidental track crossings. Figure 3.9 shows $\varepsilon_{\mathrm{total}}$ for HNL decays in the $\mu - \mu\mu$ channel for the three different MC campaigns, that correspond to the detector and pileup conditions of different data taking years and have pileup profiles as shown in fig. 2.3.

The signal DV reconstruction efficiency is observed to be robust against pileup differences whereas the effect is expected to be bigger on background since LRT fake rate increases with pileup.

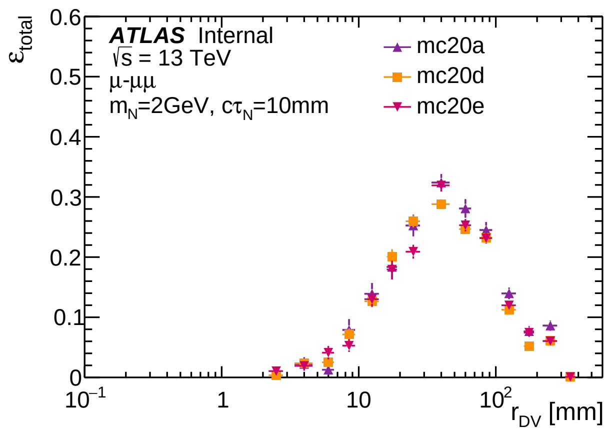

Figure 3.9: Total vertex reconstruction efficiency of $m_{\mathcal{N}} = 2$ GeV, $c\tau_{\text{HNL}} = 10$ mm HNL decays in the $\mu - \mu\mu$ channel for the mc20a, mc20d, and mc20e campaigns.

DVs with exactly two outgoing tracks are considered, where both tracks are required to be either muons or electrons, with the displaced lepton quality requirements imposed on them. The number of events with more than one such DV are found to be negligible, hence the first DV passing these conditions is selected as the HNL candidate. p_{T} cuts are applied on the DV tracks to reject soft QCD background. Finally, a fiducial selection is applied to the radial position of the DV to ensure the decay happens inside the innermost strip layer of the ID and that the decay has a minimum level of displacement from the PV. Table 3.7 summarizes the cuts applied for selection of the DV candidate.

Cut Name	Cut Description
Track multiplicity	exactly 2
Flavor	$\mu\mu$, $e\mu$, or ee
Track p_{T}	$p_{\text{T}}^{\text{lead}} > 10$ GeV, $p_{\text{T}}^{\text{sub-lead}} > 5$ GeV
Fiducial Cut	4 mm $< r_{\text{DV}} <$ 300 mm

Table 3.7: Selection criteria used for the displaced vertex.

Chapter 4

ANALYSIS STRATEGY

The goal of this analysis is to search for long-lived Heavy Neutral Leptons in the 1-20 GeV mass range. While the displaced vertex driven search provides a clean signature to work with, there are various sources of backgrounds that contaminate the targeted phase space. Section 4.1 describes these backgrounds and the kinematic cuts that remove or reduce them. Section 4.2 details the study of the background consisting of DVs with correlated tracks and its method of estimation. Section 4.3 outlines the same for DVs with uncorrelated tracks. After all the background estimation methods are in place, the signal strength is extracted using a simultaneous fit across all analysis regions as described in section 4.4.

4.1 Backgrounds

The backgrounds considered in this analysis are categorized into reducible and irreducible. While dedicated cuts are used to completely remove reducible backgrounds with negligible impact to the signal, cuts applied to suppress irreducible backgrounds also affect the signal, with the exact cut value optimized to maximize signal significance.

4.1.1 Reducible backgrounds

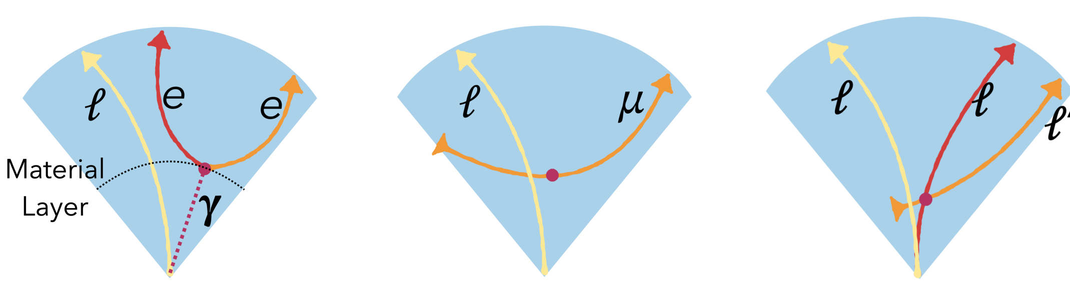

(a) $\gamma \to ee$ conversions (b) Cosmic muons (c) $Z \to \ell\ell$ decays

Figure 4.1: Illustrative topology of the reducible backgrounds in the ATLAS ID volume.

Three sources of reducible backgrounds affect this analysis, as illustrated in fig. 4.1:

1. $\gamma \rightarrow ee$ conversions: Photons interacting with the various layers of material in the ID volume recoil against the atoms and produce a e^+e^- pair. Such a ee DV, along with a prompt lepton source, contaminates the measurement phase space. To remove this background, ee DVs are required to be inside a narrower fiducial volume ($|z| < 300$ mm) and must satisfy a *material* veto, which rejects DVs which have a reconstructed position consistent with the position of the detector material. The material map, as shown in fig. 4.2, is constructed using the known positions of detector elements and from the position of low-mass vertices reconstructed in data [82]. The material veto removes 48% of the fiducial volume. The two electrons from a photon conversion are highly collinear since the recoiling nucleus takes very little momentum away from the photon. Hence, the m_{DV} distribution[1] starts at $\sqrt{2}m_e$ and falls sharply. Since the m_{DV} calculation makes a pion mass assumption on the outgoing tracks, the lower mass bound is given by $\sqrt{2}m_{\pi^+} \approx 280$ MeV. Hence, to further reject any remaining conversions, a $m_{\mathrm{DV}} >$ 500 MeV cut is applied on ee DVs. The assumption on the mass does not affect the impact of the cut since an m_{DV} measurement of the scale $\sqrt{2}m_e$ cannot be resolved by ATLAS. The material veto also hurts the signal efficiency since it removes almost half the fiducial volume from the analysis. However, in the current context, the photon conversion background is still classified as a reducible background since it is completely removed after the above cuts.

2. Cosmic muons: Even though ATLAS is 100 m below the ground, the open shaft above it allows a large flux of muons from the the interaction of cosmic rays with the atmosphere to enter the detector. One such cosmic muon passing through the ATLAS ID can be incorrectly reconstructed as two tracks, with them together reconstructed as the two legs of one secondary vertex. Such a DV is characterized by *back-to-back* tracks i.e. $\eta_{\mathrm{track1}} \approx -\eta_{\mathrm{track2}}$ and $\phi_{\mathrm{track1}} - \phi_{\mathrm{track2}} \approx \pi$. Hence, a cut on the cosmic

[1] $m_{\mathrm{DV}} \approx \sqrt{E_1 \cdot E_2(1-\cos\alpha) + 2m_1m_2}$ for relativistic objects with energies and masses E_1, E_2 and m_1, m_2, respectively.

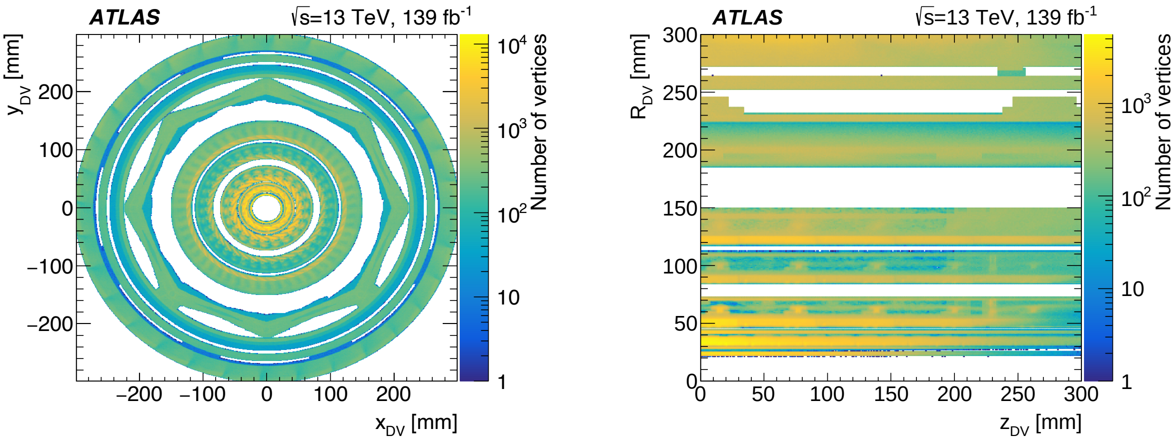

(a) Positions in $x - y$ plane

(b) Positions in $r - z$ plane

Figure 4.2: The positions of reconstructed displaced vertices that are vetoed by the material map. [82]

separation variable, $\sqrt{(\Sigma_{\text{DV tracks}}\eta)^2 + (\pi - |\Delta_{\text{DV tracks}}\phi|)^2} > .05$, is applied to reject this background. The exact cut value was optimized for the 2022 analysis, and was found to be sufficient for rejecting all background without affecting signal.

3. $Z \to \ell\ell$ decays: One of the leptons from a Z decay can be paired with a third uncorrelated lepton to give a lepton+DV topology probed in this analysis. To reject this background, a veto on the $m_{\ell\ell}$ variable around the Z mass, $|m_{\ell\ell} - 90 \text{ GeV}| > 10$ GeV is applied to same-flavor opposite-sign (SFOS) pairs of leptons, where one is a prompt lepton and the other is a DV lepton.

Table 4.1 summarize the cuts applied to reject the three reducible backgrounds that affect this analysis.

Cut Name	Cut Description
Photon Conversions Veto	Material veto applied to ee DVs
Low Mass DV Veto	$m_{\text{DV}} > 500$ MeV for ee DVs
Cosmic Muon Veto	$\sqrt{(\Sigma_{\text{DV tracks}}\eta)^2 + (\pi - \|\Delta_{\text{DV tracks}}\phi\|)^2} > .05$
$Z \to \ell\ell$ Veto	$\|m_{\ell_{\text{DV}}\ell_{\text{prompt}}} - 90 \text{ GeV}\| > 10$ GeV for SFOS pairs

Table 4.1: Analysis cuts applied to reject reducible background.

4.1.2 Irreducible backgrounds

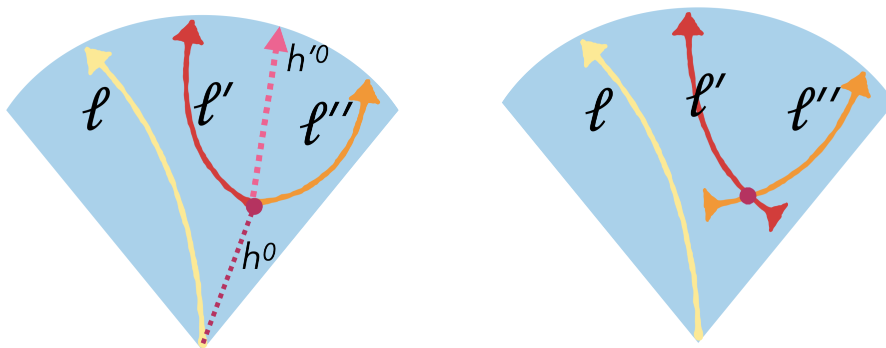

(a) Correlated background (b) Uncorrelated background

Figure 4.3: Illustrative topology of the irreducible backgrounds in the ATLAS ID volume.

Two sources of irreducible backgrounds affect this analysis, as illustrated in fig. 4.3:

1. Correlated background: The largest background that affects this analysis is the lepton+DV signature coming from correlated decay sources. Heavy hadrons, such as b-hadrons and secondary J/ψ, mimic signatures from BSM LLPs and decay macroscopically in the ID volume. While these decays are often multi-legged, two of its legs can be captured by VSI and hence contaminate the phase space targeted in this analysis. Section 4.2 describes the detailed study of this background and the method of its estimation.

2. Uncorrelated background: A small background contribution comes in the form of uncorrelated track pairs, or *random crossings* (RC), that happen to cross in an event and then captured by VSI. The DV is not a genuine decay but rather an accidental point of crossing. Section 4.3 describes the method used in the analysis to estimate the contribution from random crossing background.

4.2 Correlated Background

As previously stated, the biggest background in this analysis is the background from the production of heavy flavor (HF) hadrons which decay to leptons that are displaced with respect to the primary vertex. These are predominantly found in the $m_{\mathrm{DV}} < 5$ GeV phase space. Such a phase space was completely removed from consideration in the 2022 analysis due to the lack of a robust method to estimate them which significantly reduced sensitivity to low mass HNLs. In this analysis, MC simulations of SM processes (specifically of the $t\bar{t}$ and V+jets processes) are used to understand and estimate this background. A signal rich phase space, called the Signal Region (SR), is defined to look for HNLs in a wide lifetime-coupling parameter space. A non-overlapping phase space rich in correlated background, called the Heavy Flavor Control Region (HF-CR), is defined to obtain the normalization of this background in a data-driven manner and reduce the uncertainties on the estimate.

4.2.1 Analysis regions

All events considered in this analysis are required to have at least one primary vertex with two tracks with $p_{\mathrm{T}} > 500$ MeV. The hard-scatter vertex is selected among all reconstructed primary vertices as the one with the largest Σp_{T}^2, where the sum is over all the tracks attached to the PV.

A pre-selection cut is applied for the selection of signatures from a W-decay. The tri-lepton invariant mass ($m_{\ell\ell\ell}$) is used as a proxy for the W mass. Events are required to have 40 GeV$< m_{\ell\ell\ell} <$90 GeV, where the lower cut is more liberal compared to the upper cut to account for energy losses from a neutrino emission. A J/ψ-veto is applied for ee and $\mu\mu$ DVs requiring $|m_{\mathrm{DV}} - 3.1\ \mathrm{GeV}| > 0.1$ GeV.

The final pre-selection cut is applied using a discriminating variable optimized to significantly suppress the contribution from hadron decays. The variable $\mathcal{S}$ is defined as

$$\mathcal{S}^2 = \mathbf{\Delta}_{\mathrm{3D}}(\mathrm{PV,DV})^T \cdot \mathbf{Cov}(\Delta_{\mathrm{3D(PV,DV)}})^{-1} \cdot \mathbf{\Delta}_{\mathrm{3D}}(\mathrm{PV,DV}), \tag{4.1}$$

where $\mathbf{\Delta}_{\mathrm{3D}}(\mathrm{PV,DV})$ is the 3D vector of the distance between the primary- and secondary-vertex position, and $\mathbf{Cov}(\Delta_{\mathrm{3D(PV,DV)}})$ is the covariance matrix of the difference between the primary- and secondary-vertex position. $\mathcal{S}$ can be interpreted as a significance of the PV-DV distance. Roughly speaking, DVs with lower values of $\mathcal{S}$ have larger $p_{\mathrm{T}}^{\mathrm{DV\ tracks}}$ and lower r_{DV} and ΔR between the DV tracks.

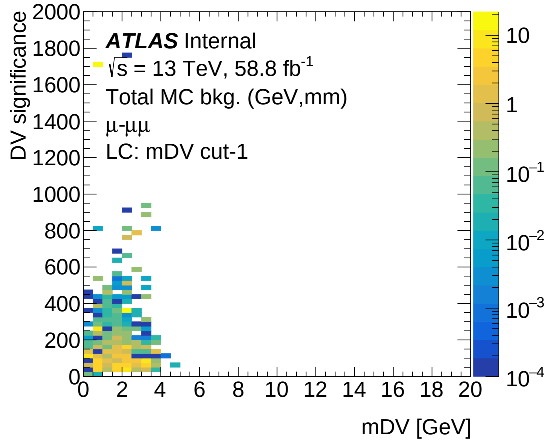

(a) Hadron decay background

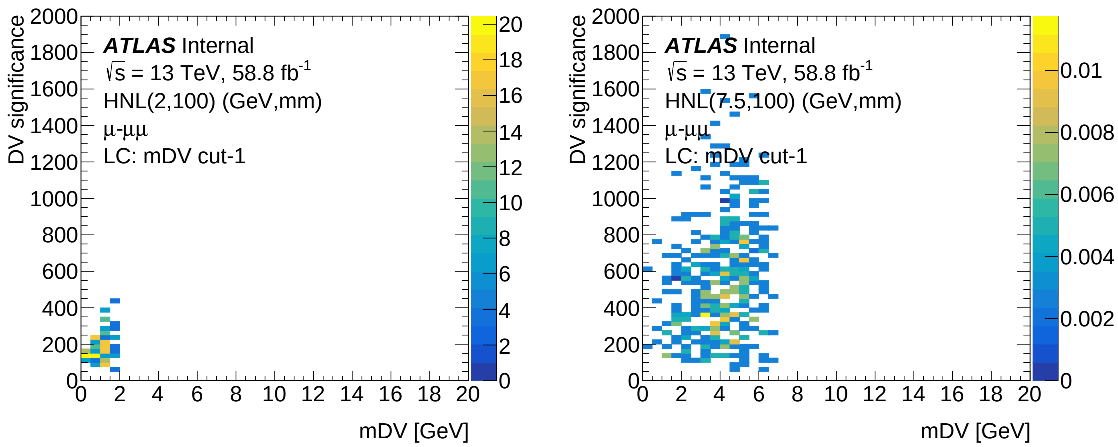

(b) m_{HNL} =2 GeV, $c\tau_{\mathrm{HNL}}$=1000 mm HNL signal(c) m_{HNL} =7.5 GeV, $c\tau_{\mathrm{HNL}}$=100 mm HNL signal

Figure 4.4: Distribution of events as a function of the $\mathcal{S}$ variable, referred to as 'DV significance' in the plots and m_{DV}.

The distribution of events as a function of $\mathcal{S}$ and the mass of the displaced vertex, m_{DV} is shown in fig. 4.4 for the hadron decay background predicted by the combination of $t\bar{t}$ and V+jets MC simulations as well as for representative signals in for pairs $(m_{\mathrm{HNL}}, c\tau_{\mathrm{HNL}})$ = (2 GeV, 1000 mm) and (7.5 GeV, 100 mm).

From fig. 4.4, it is clear that the hadron decay backgrounds are characterised by lower values of $\mathcal{S}$, while for the signal samples the discriminant variable peaks at values above 100. Furthermore, the events from the SM processes are present only for displaced vertices masses below 5 GeV, as they are mostly produced in B-hadron decay chains. For these reasons, I decided to apply the following selection to all the events in the analysis: $\mathcal{S} > 100$ *if* $m_{\mathrm{DV}} < 5$ GeV. This means that all the events outside of the rectangle defined by the four coordinates (0,0), (5,0), (0,100), (5,100) in the (m_{DV} [GeV], $\mathcal{S}$) plane are retained.

The SR is defined by requiring, on top of the events pre-selection, that the two displaced leptons satisfy an isolation criterion, using the `Loose_VarRad` [78] WP for electrons, and the `PflowLoose_VarRad` WP [55] for muons. Leptons from the DV tend to be collimated and hence close-by-effects[2] that might appear in the evaluation of isolation variables are taken into account and correctly subtracted from both the track term and the calorimeter term that define the isolation working points. This procedure is done since Run-1 in $H \to ZZ^* \to 4\ell$ analyses, and generally applies to boosted multi-lepton signatures utilizing lepton isolation. The HF-CR, on the other hand, is defined adding a requirement of at least one non-isolated displaced lepton. To further suppress signal contamination in this region, at least one b-tagged jet is required. This requirement does not bias the kinematics of the correlated background in the HF-CR away from the SR. Lastly, since this analysis probes HNLs with masses up to 20 GeV, the phase space beyond $m_{\mathrm{HNL}} > 20$ GeV is not considered in the SR. Table 4.2 summarizes the cuts used for the SR and the HF-CR. There are 6 SRs and 6 HF-CRs defined in this analysis, one each per channel.

Figures 4.5 to 4.10 show representative kinematics of important properties of the DV in the event and the tracks in it, the prompt lepton, and the $\mathcal{S}$ discriminant variable for events in the HF-CR for all the six channels.

[2]The calculation of isolation variables uses low-level information from the tracker and the calorimeter without relying on object-level reconstruction. When two leptons are close to each other, they will contribute to the calculation of each other's isolation variables. However, since isolation is only supposed to identify leptons in jets, such a lepton-close-by-effect is an artefact of the method and needs to be corrected.

Cut Name	Cut Description	
	Object Selection	
Primary Vertex	Largest $\Sigma_{\text{tracks}} p_{\text{T}}^2$	
Prompt Lepton	see table 3.2	
Trigger	see table 3.3	
Displaced Lepton	see tables 3.4 and 3.5	
Overlap Removal	see table 3.6	
Displaced Vertex	see table 3.7	
	Pre-selection	
$\gamma \to ee$ Veto	Material veto, $\lvert z \rvert < 300$ mm and $m_{\text{DV}} > 500$ MeV for ee DVs	
Cosmic Muon Veto	$\sqrt{(\Sigma_{\text{DV tracks}}\eta)^2 + (\pi - \lvert\Delta_{\text{DV tracks}}\phi\rvert)^2} > .05$	
$Z \to \ell\ell$ Veto	$\lvert m_{\ell_{\text{DV}}\ell_{\text{prompt}}} - 90 \text{ GeV}\rvert > 10$ GeV for SFOS pairs	
W-decay Selection	40 GeV $< m_{\ell\ell\ell} <$ 90 GeV	
J/ψ Veto	$\lvert m_{\text{DV}} - 3.1 \text{ GeV}\rvert > 0.1$ GeV for ee and $\mu\mu$ DVs	
Discriminant	$\mathcal{S} > 100$ *if* $m_{\text{DV}} < 5$ GeV	
	Region Specific Cuts	
	SR	**HF-CR**
DV Lepton Isolation	Both isolated	≥ 1 non-isolated
b-jets	-	≥ 1
m_{HNL} Limit	< 20 GeV	-

Table 4.2: Object selection, pre-selection, and other additional cuts used to define the non-overlapping Signal Region and the Heavy Flavor Control Region.

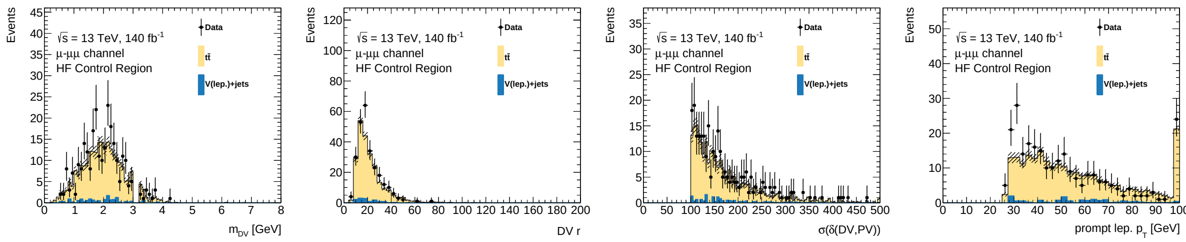

(a) m_{DV} [GeV] (b) r_{DV} [mm] (c) $\mathcal{S}$ (d) prompt lep. p_{T} [GeV]

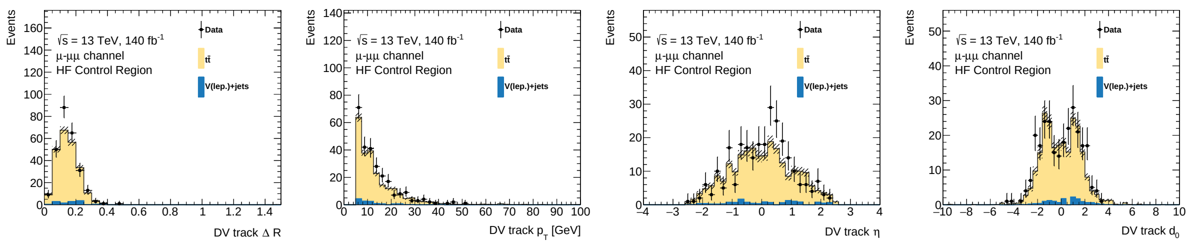

(e) ΔR DV tracks (f) p_{T} of DV tracks [GeV] (g) η of DV tracks (h) d_0 of DV tracks [mm]

Figure 4.5: Representative kinematics of the displaced vertex, tracks in the displaced vertex, and of the prompt lepton measured in ATLAS data and modeled by MC simulations in the $\mu - \mu\mu$ Heavy Flavor Control Region.

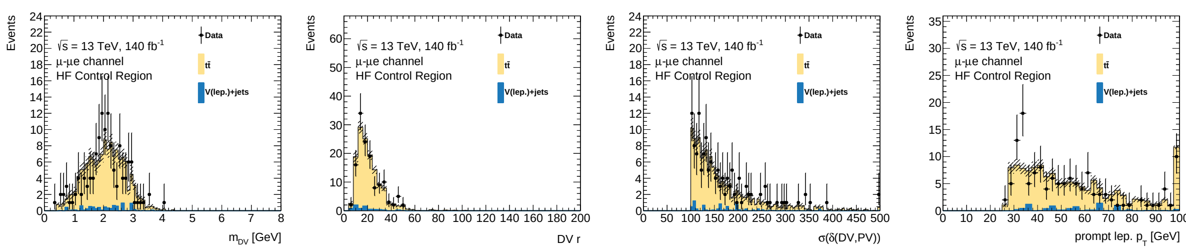

(a) m_{DV} [GeV] (b) r_{DV} [mm] (c) $\mathcal{S}$ (d) prompt lep. p_{T} [GeV]

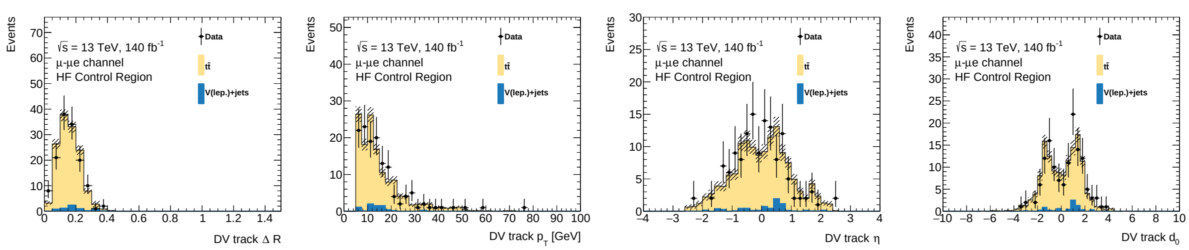

(e) ΔR DV tracks (f) p_{T} of DV tracks [GeV] (g) η of DV tracks (h) d_0 of DV tracks [mm]

Figure 4.6: Representative kinematics of the displaced vertex, tracks in the displaced vertex, and of the prompt lepton measured in ATLAS data and modeled by MC simulations in the $\mu - \mu e$ Heavy Flavor Control Region.

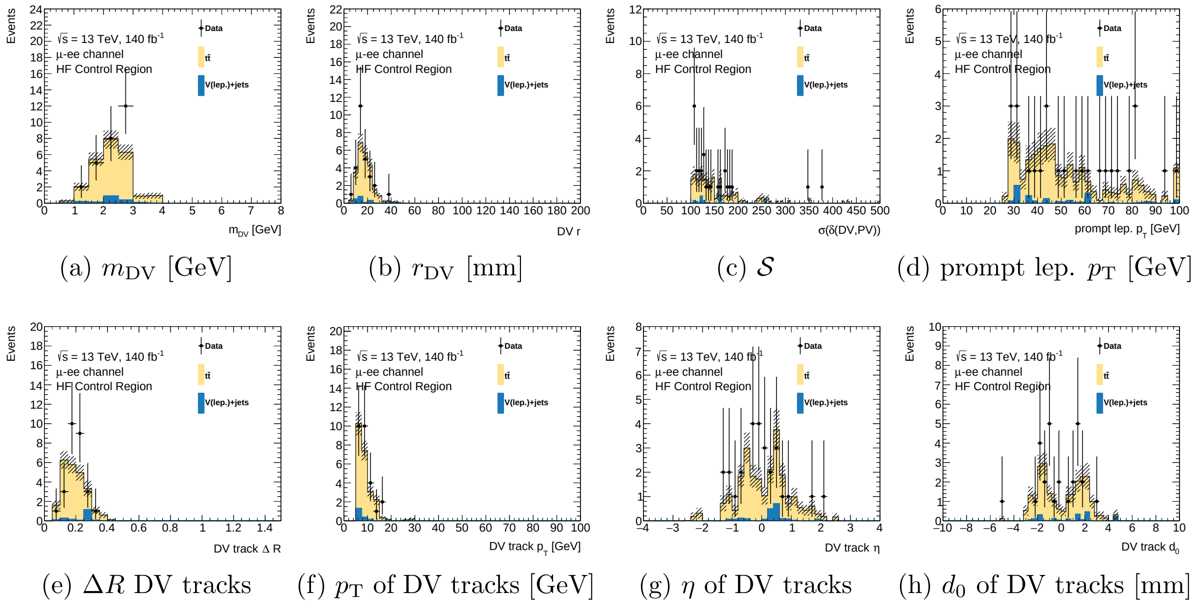

(a) m_{DV} [GeV] (b) r_{DV} [mm] (c) $\mathcal{S}$ (d) prompt lep. p_{T} [GeV]

(e) ΔR DV tracks (f) p_{T} of DV tracks [GeV] (g) η of DV tracks (h) d_0 of DV tracks [mm]

Figure 4.7: Representative kinematics of the displaced vertex, tracks in the displaced vertex, and of the prompt lepton measured in ATLAS data and modeled by MC simulations in the $\mu - ee$ Heavy Flavor Control Region.

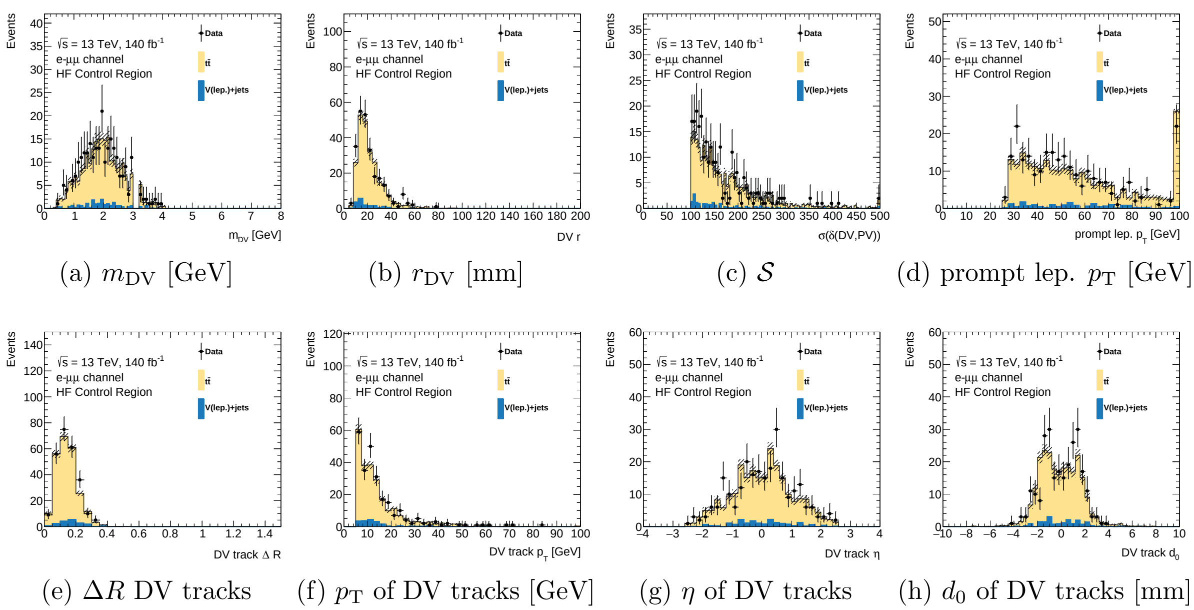

(a) m_{DV} [GeV] (b) r_{DV} [mm] (c) $\mathcal{S}$ (d) prompt lep. p_{T} [GeV]

(e) ΔR DV tracks (f) p_{T} of DV tracks [GeV] (g) η of DV tracks (h) d_0 of DV tracks [mm]

Figure 4.8: Representative kinematics of the displaced vertex, tracks in the displaced vertex, and of the prompt lepton measured in ATLAS data and modeled by MC simulations in the $e - \mu\mu$ Heavy Flavor Control Region.

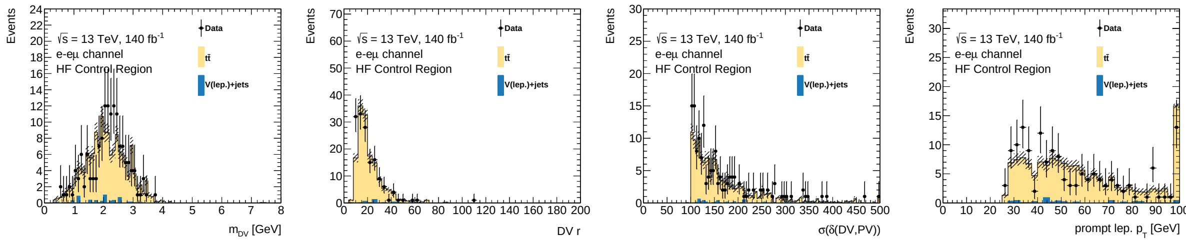

(a) m_{DV} [GeV] (b) r_{DV} [mm] (c) $\mathcal{S}$ (d) prompt lep. p_{T} [GeV]

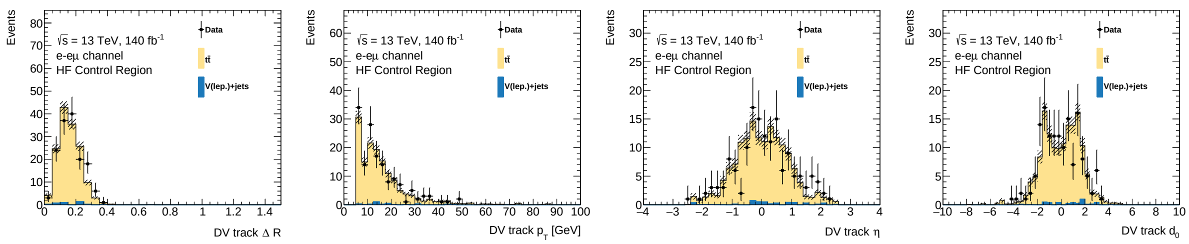

(e) ΔR DV tracks (f) p_{T} of DV tracks [GeV] (g) η of DV tracks (h) d_0 of DV tracks [mm]

Figure 4.9: Representative kinematics of the displaced vertex, tracks in the displaced vertex, and of the prompt lepton measured in ATLAS data and modeled by MC simulations in the $e - e\mu$ Heavy Flavor Control Region.

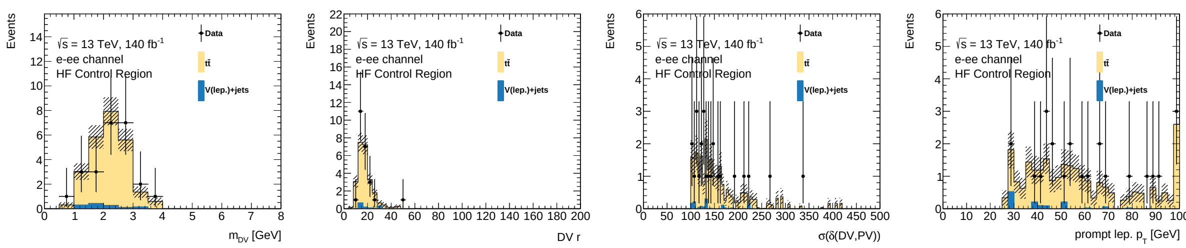

(a) m_{DV} [GeV] (b) r_{DV} [mm] (c) $\mathcal{S}$ (d) prompt lep. p_{T} [GeV]

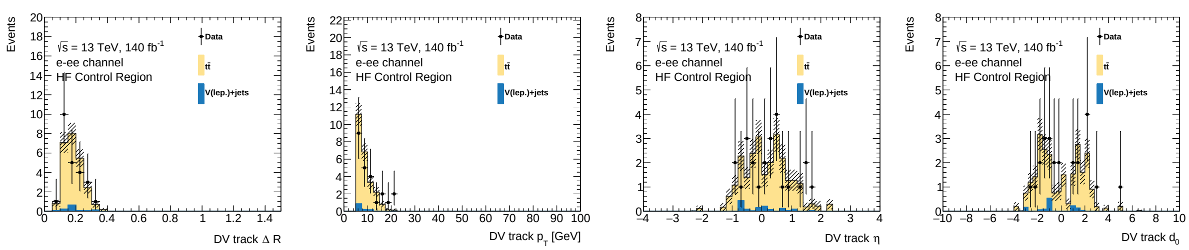

(e) ΔR DV tracks (f) p_{T} of DV tracks [GeV] (g) η of DV tracks (h) d_0 of DV tracks [mm]

Figure 4.10: Representative kinematics of the displaced vertex, tracks in the displaced vertex, and of the prompt lepton measured in ATLAS data and modeled by MC simulations in the $e - ee$ Heavy Flavor Control Region.

4.2.2 Truth classification

As previously mentioned, the kinematics of the hadron decay background are modeled using MC simulations of $t\bar{t}$ and V+jets processes and the normalization is obtained from data. Other processes, such as di-boson, multi-jet, $b\bar{b}$, which could lead to the considered final state were studied, and the contribution from them was found to be negligible. Such simulations offer a way to study the physical source of the final state using the *truth origin*[3]. of the tracks in the DV. The truth origin is calculated by looking at the parent information of the truth-matched particle of an ID track. Based on the kind of the parent, the truth origin is delineated into a set of categories.

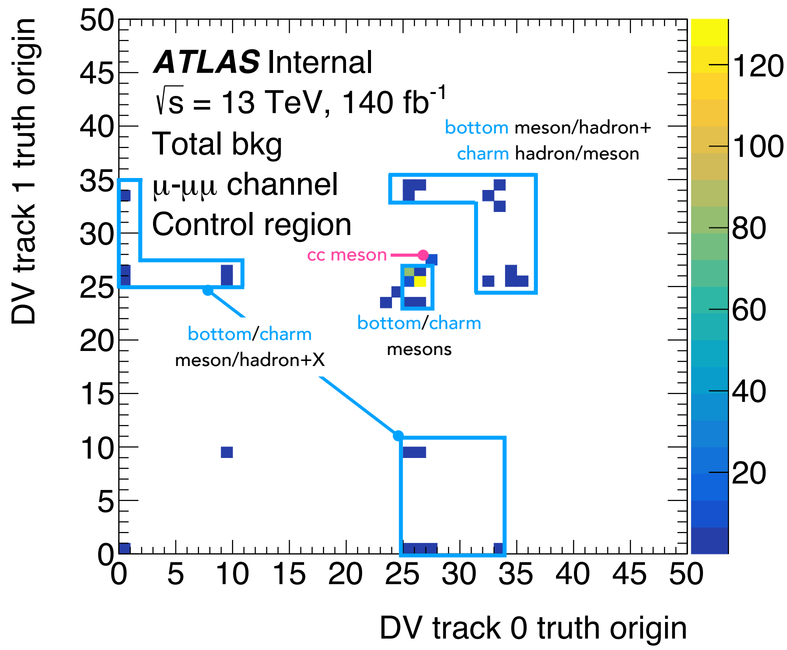

Figure 4.11: Event categories based on the truth classification of tracks associated to the displaced vertex in the control region of the $\mu - \mu\mu$ channel for the total MC background.

In fig. 4.11, an example plot shows the grouping of different kind of DVs shown for the events passing the HF-CR selection of the $\mu - \mu\mu$ channel. This region has been chosen

[3]The truth origin of a track, an integer value ranging from 0 to 45, demarcates the nature of the particle which decayed to give the DV track using the `MCTruthClassifier` [83] package. The available categories can be seen in the ATLAS reconstruction framework `https://gitlab.cern.ch/atlas/athena/-/blob/24.0/PhysicsAnalysis/MCTruthClassifier/MCTruthClassifier/MCTruthClassifierDefs.h`.

because it is a high-statistics region and it allows to clearly highlight the different contributions. The truth origin of one of the two tracks in the DV is plotted on the x-axis and the truth origin of the other one is plotted on the y-axis. As can be seen from the z-axis scale, the two main contributions to the total background are events with truth origins for (track 1, track 2) = (25,26) and (26, 25): these correspond to tracks associated to the decay productions of B- (26) and D-mesons (25). This is because B mesons can decay into D mesons plus a charged lepton. Subsequently, real leptons can also be produced in multiple steps of the D-meson decay-chain. Such cases are reconstructed as a single DV with two charged leptons, thus passing the selection of the analysis. The presence of additional (low quality) tracks in the decay chain are less likely to be captured by the vertexing algorithm, since the track attachment step is modified to only use tracks that pass the track seeding requirements.

Apart from the main source of HF background, some events are also found with one of the tracks having truth origin 23 or 24, corresponding to leptons originating from light hadrons or strange mesons, respectively. These can be found in the decay chains of B and D mesons as well.

Another non negligible component comes from pairs of track with true origin equal to 27: these are tracks associated to the production of $c\bar{c}$ mesons, i.e. J/ψ contribution. This is suppressed by the J/ψ veto introduced in this round of the analysis, however it still contributes and is localised in the m_{DV} distribution around 3.1 GeV.

Additional background is found to be associated to the production of heavy flavour baryons, both containing b- and c- quarks. This population of events leads to DVs in which the tracks are classified as coming from bottom meson + charm hadron or charm meson + bottom hadron.

The last population, significantly smaller than the ones above, consist of events in which at least one the tracks associated to the displaced vertex comes from an unknown (0) origin. This reflects the presence of real B- or D-meson production which undergo subsequent decays into leptons, with the other lepton comes from either jets in the events located close to the lepton, or, less likely, directly from the underlying event or a pileup track. This category of events is negligible with respect to the others.

A very minor contribution comes from two tracks with truth origin 9, corresponding to tracks associated to the decay of τ leptons. This component is also negligible and hence ignored in the rest of the analysis.

The population of events at (0,0) is instead cases in which two uncorrelated tracks crossing and producing a secondary vertex. These events will not be taken from the MC simulated samples. Instead, a dedicated estimate of this background is performed, as is outlined in section 4.3.

Given the grouping of the events discussed above, the following two categories are defined:

1. Heavy Flavour hadron production: all the events that have as truth origin of at least one track either b- or c-meson or baryon and the events that have the 27 as the truth origin of both tracks, which corresponds to $c\bar{c}$ meson production,

2. *other*: rest of the events. This category has a minor contribution and has been considered only for *in itinere* data/MC plots.

The residual contribution of the *other* component, **not included in the fit templates**, is expected to be mostly covered by the random crossing background estimate, and eventually absorbed by a global MC normalisation factor which will be used to correct the prediction of the HF background component using data in the dedicated control regions.

Figure 4.12 shows the truth origin distribution for the total MC prediction for the SRs in the six channels. The background composition is very similar to that in the $\mu - \mu\mu$ HF-CR, with the production of heavy-flavour hadrons being the leading contribution. One notable difference between the $\mu - \mu\mu$ HF-CR, discussed previously, and the SR, discussed here, is that there are almost no events with tracks with unknown truth origin (0). This is related to the fact that in the signal regions the two reconstructed leptons matched to the tracks

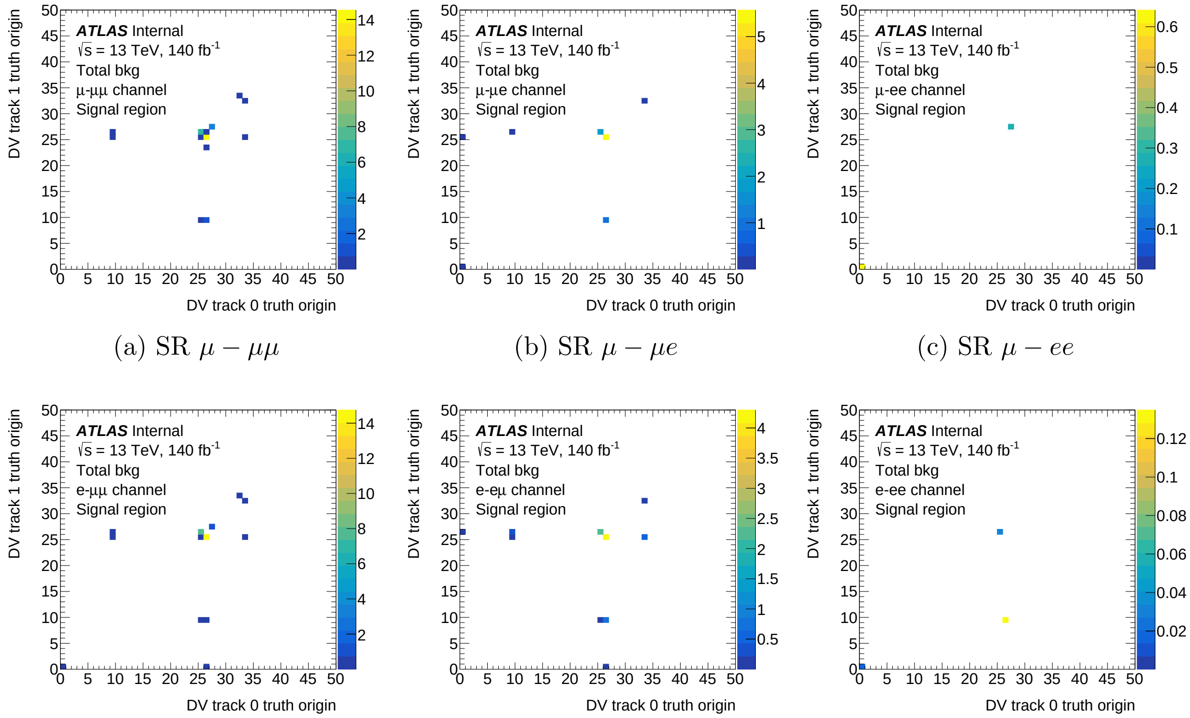

(a) SR $\mu - \mu\mu$ (b) SR $\mu - \mu e$ (c) SR $\mu - ee$

(d) SR $e - \mu\mu$ (e) SR $e - e\mu$ (f) SR $e - ee$

Figure 4.12: Truth origin distributions of track pairs associated to displaced vertices in the six leptonic signal region channels for the total MC background.

associated to the DV are required to be isolated. Such a requirement reduces the probability that one real lepton produced in a DV is associated to a random track coming from a jet close-by.

4.2.3 Background prediction

To exploit the fact that the different background components have different DV mass distributions, the chosen variable in the HF-CRs for the statistical analysis, described later in section 4.4, is m_{DV} itself. On the other hand, since the HNL signal extraction is done from the SRs, the reconstructed HNL mass (m_{HNL}) introduced in the 2022 analysis (details of calculation in chapter B of the appendix) is used for its strong discrimination power. The reconstruction uses energy-momentum conservation in the HNL production ($W \to \mathcal{N}\ell_1$) and decay ($\mathcal{N} \to \ell_1\ell_2\nu$) to get around the lack of measurement of the neutrino kinematics.

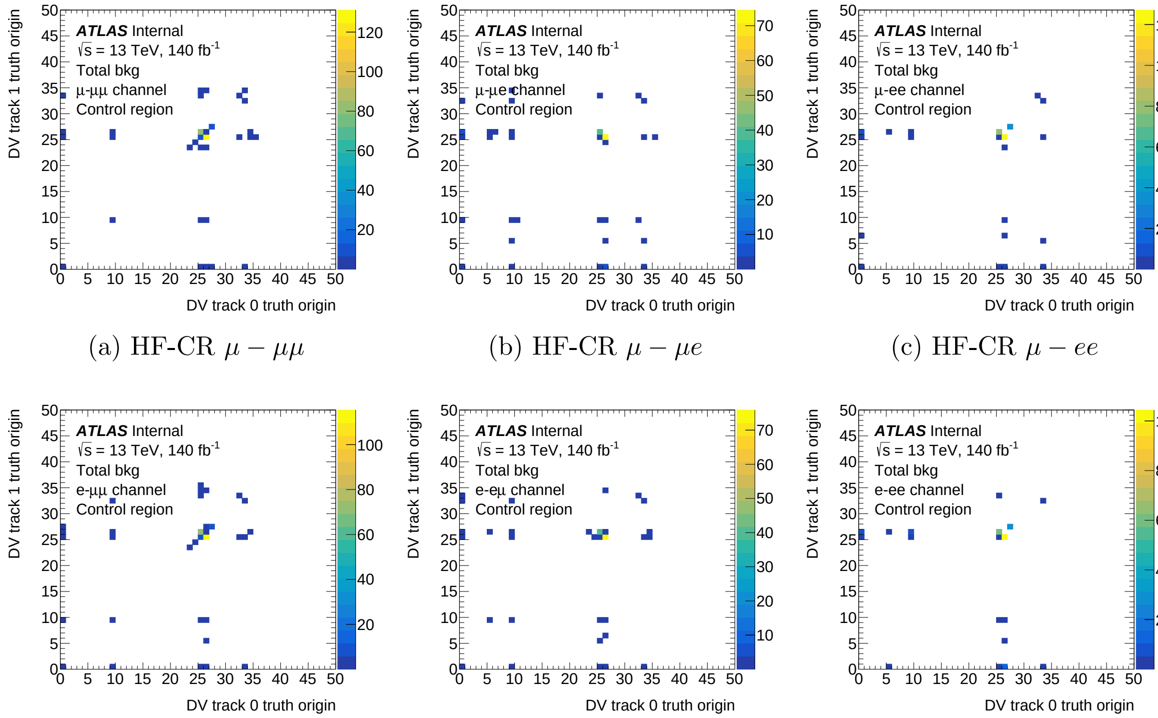

(a) HF-CR $\mu - \mu\mu$ (b) HF-CR $\mu - \mu e$ (c) HF-CR $\mu - ee$

(d) HF-CR $e - \mu\mu$ (e) HF-CR $e - e\mu$ (f) HF-CR $e - ee$

Figure 4.13: Truth origin distributions of track pairs associated to displaced vertices in the six leptonic control region channels for the total MC background.

The MC samples are used to build the templates of predictions used in the statistical analysis. Figure 4.14 shows the reconstructed m_{HNL} for signal samples with various simulated masses and $c\tau_{\mathrm{HNL}} = 10$ mm. The m_{HNL} resolution for channels with $\mu\mu$ DVs is the best, followed by μe and ee DVs. This is due to the use of the kinematics of the track object and not of the lepton object in the reconstruction of this variable. Since GSF tracks don't fully account for radiation losses from electrons, they give a worse momentum measurement of the HNL decay products, and hence a worse reconstructed m_{HNL} value.

The available statistics from MC simulations is sufficient in the HF-CRs, therefore, the expected background MC events are used. A systematic uncertainty equal to the relative statistical uncertainty of the background MC predictions is assigned to this template. In the HF predictions in the SR, instead, some un-physical statistical fluctuations are observed. Furthermore, insufficient statistics in the SR from the MC samples lead to no background

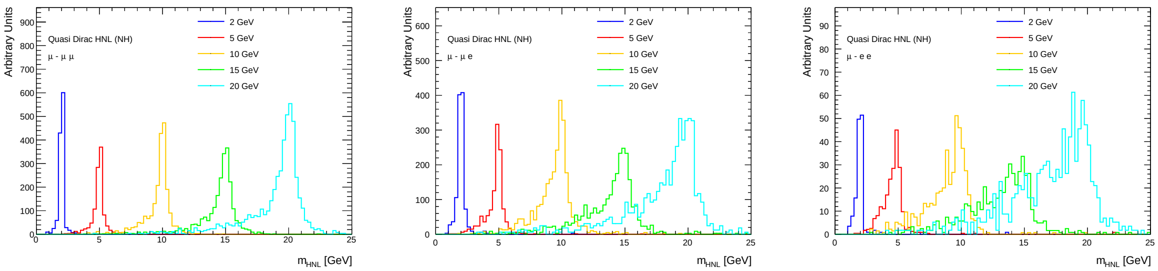

(a) $\mu - \mu\mu$ (b) $\mu - \mu e$ (c) $\mu - ee$

Figure 4.14: Reconstructed m_{HNL} distributions for signal samples with $c\tau_{\text{HNL}} = 10$ mm. The signals are scaled to arbitrary values for side-by-sde comparison between different mass points.

prediction beyond $m_{\text{HNL}} = 7$ GeV. In order to reduce the impact of such artefacts, the shape of the m_{HNL} distribution of the HF background in the SR is obtained from a shape region (ShR) with requirements on the DV lepton isolation and $m_{\ell\ell\ell}$ removed. Since similar kinematics and compositions are observed in all the regions for this background, events from the ShR of all six channels are combined to obtain the shape. Finally, the shape is scaled back to the event yield predicted by MC simulations in the six individual channels. A dedicated systematic uncertainty is assigned to this method to cover the assumptions made with the details mentioned in section 4.4.

Figure 4.15 shows the template building procedure for the HF background component in the SR. It shows the merging of the six different channels (bottom row), first in $x - \mu\mu$, $x - \mu e$ and $x - ee$ channels (middle row), and then all the channels merged (top plot). The yield from ShR in each case is shown, normalized to the nominal MC event yield for a one-to-one shape comparison. The original MC at each level is shown with black dots and the shape from ShR normalized to the original MC in individual channels in shown in dashed red lines. For the sake of illustration, the ShR shape at the DV flavor level is also shown in blue, however this shape is not used directly in the analysis. Finally, the shape from the ShR for all channels combined, which is used as the SR background estimate shape in the SR, is shown by the green line. The agreement between the original MC predictions and the

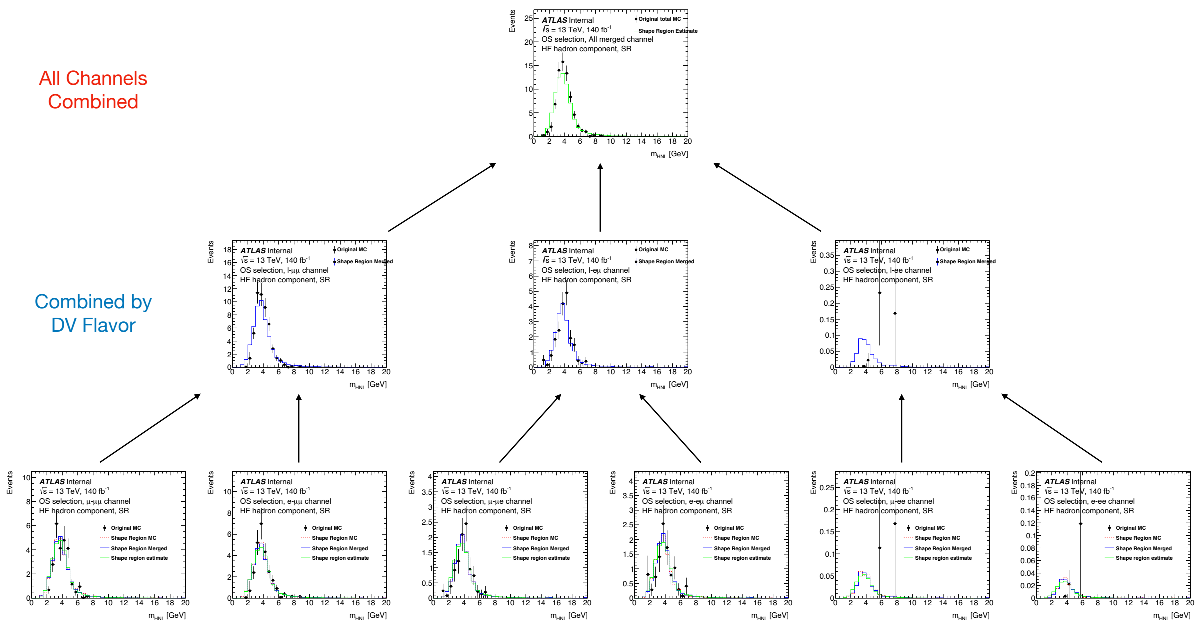

Figure 4.15: Template building method used for the HF background in the Signal Regions.

ShR shape is sufficiently good across all channels in the $m_{\text{HNL}} = 1 - 7$ GeV range, with the ShR prediction also giving a non-zero prediction in the tails of the distribution, i.e. in the $m_{\text{HNL}} = 7 - 20$ GeV range. The templates obtained with this method are shown in fig. 4.16.

4.3 Uncorrelated Background

The Random Crossing background consists of a pair of uncorrelated tracks crossing each other in space and accidentally getting reconstructed as a DV paired with a prompt lepton to yield an event in our measured phase space. These exact characteristics of the RC background are used to estimate them, the method for which was developed for [84]. Figure 4.17 illustrates the steps used to create RC DV candidates in a data driven manner.

In the first step, a pool of tracks passing the track selection requirements of VSI is created from all available data events. From this pool, track pairs are randomly selected (with replacement) which are required to come from different events. In the last step, these track pairs are passed through VSI and a DV reconstruction is attempted. The result is

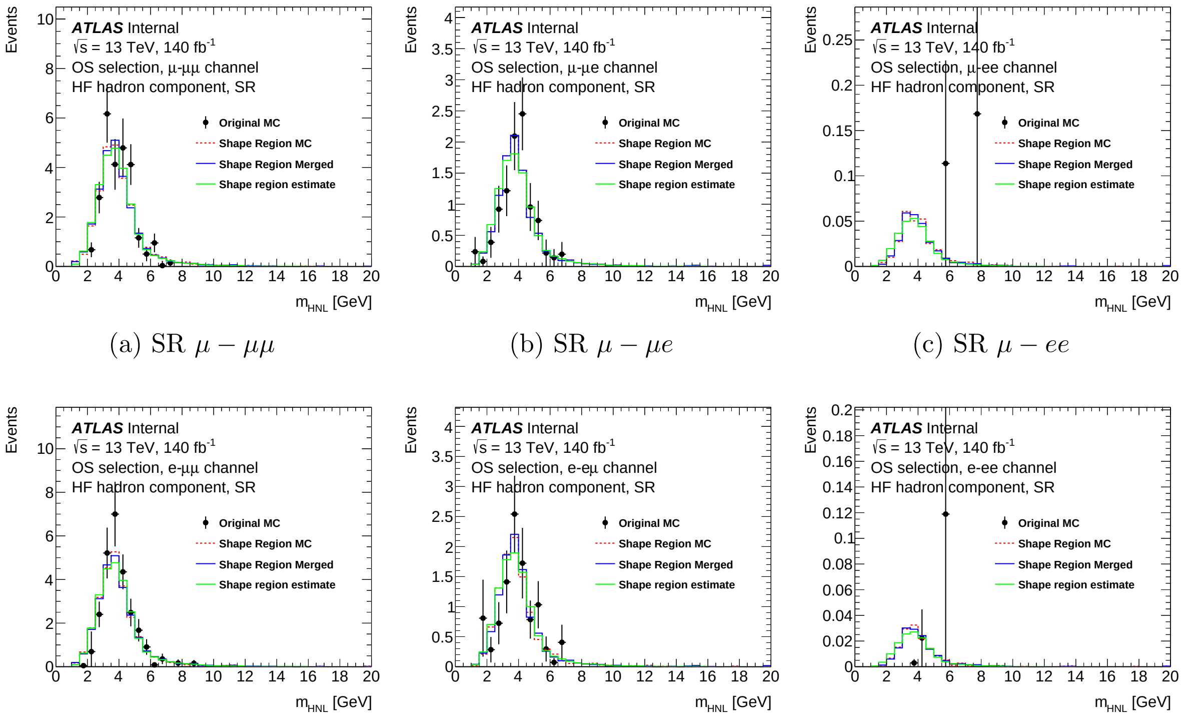

(a) SR $\mu - \mu\mu$ (b) SR $\mu - \mu e$ (c) SR $\mu - ee$

(d) SR $e - \mu\mu$ (e) SR $e - e\mu$ (f) SR $e - ee$

Figure 4.16: Comparison of the original MC predictions and different smoothed predictions for the HF background in the Signal Regions.

a set of artificial DVs representative of the vertices expected to result from random track crossings.

The probability to obtain a DV from a track pair (p(DV—pair)) is given by:

$$p(\text{DV—pair}) \approx \lim_{N_{\text{Attempts}} \to \infty} \frac{N_{\text{DV}}}{N_{\text{Attempts}}}. \tag{4.2}$$

where N_{Attempts} is the number of times a DV reconstruction is attempted using track pairs and N_{DV} is the number of times a DV was successfully reconstructed in the attempts. The number of predicted RC DVs ($N_{\text{Predicted crossings}}$), hence, follows as:

$$N_{\text{Predicted crossings}} = p(\text{DV}|\text{Pair}) \times N_{\text{Real}} = N_{DV} \times \frac{N_{\text{Real}}}{N_{\text{Attempts}}}. \tag{4.3}$$

Therefore, each artificially reconstructed RC DV is weighted by a factor $\frac{N_{\text{Real}}}{N_{\text{Attempts}}}$, where N_{Real}

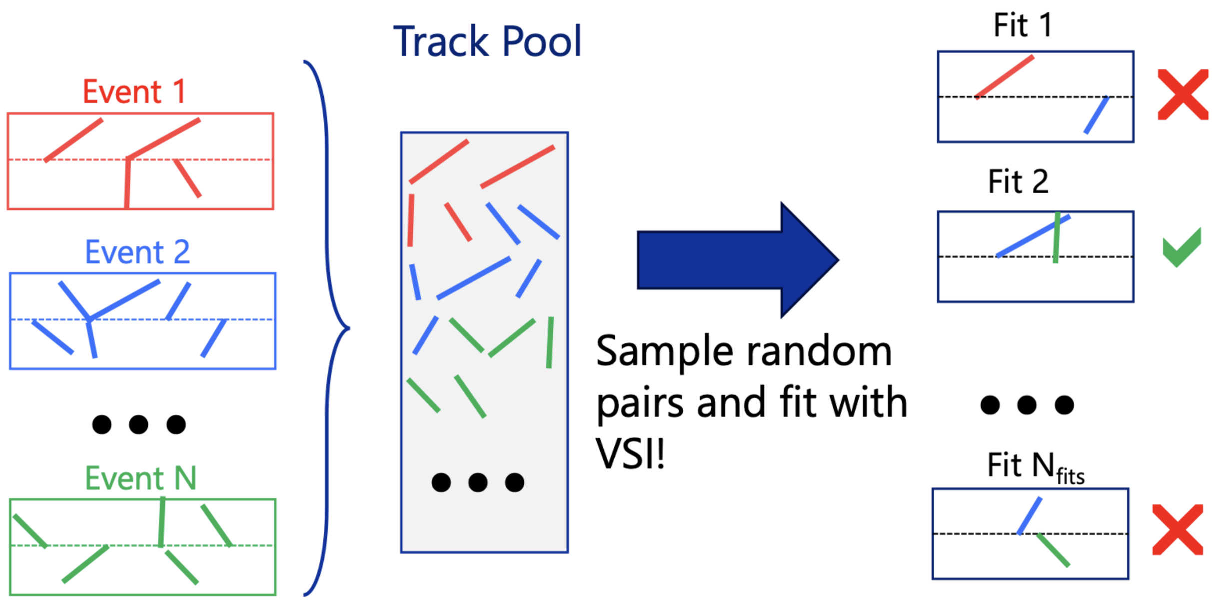

Figure 4.17: Schematic illustration of track pool creation from different events used to create randomly paired tracks for artificial vertex creation to estimate the random track crossing background.

is the number of combinatorial track pairs found in data. Care is taken to ensure that eq. (4.3) are representative of unrelated track pairs. Correlated track pairs from decaying resonances should not enter N_{Real}, as such pairs do not occur randomly, and the DV that may result from them are estimated using dedicated other techniques. To ensure this, both N_{Real} and N_{Attempts} are obtained by counting pairs with $m_{\text{Pair}} > 6$ GeV to reject low-mass hadron decays, and $|\Delta\phi_{\text{Pair}}| > 2.4$ to reject Drell-Yan and resonant Z decay events. In addition, the normalisation is performed separately for events that contain prompt leptons and events that don't. The normalisation weights are obtained separately for each DV flavor.

Finally, the RC DV candidate is combined with the prompt lepton (which is required to come with one of the tracks in the pool) to yield a complete signature defining the RC background in this search. The kinematics of this 3-lepton system are then used to reconstruct variables that this analysis is sensitive to, such as m_{DV} and m_{HNL}.

The kinematics and normalization of the RC estimate are successfully validated in a phase space rich in this background. The $m_{\ell\ell\ell}$ pre-selection is removed, the fiducial volume cut is relaxed (2 mm $< r_{\text{DV}} <$ 300 mm, and an added high m_{DV} ($>$ 2 GeV) requirement is imposed. From this phase space, all events passing the SR cuts are veto-ed to define the validation region. The background is found to be extremely suppressed with very low expected yields in all regions, as shown in table 4.3.

Region	$\mu-\mu\mu$	$\mu-\mu e$	$\mu-ee$	$e-\mu\mu$	$e-e\mu$	$e-ee$	Total
SR	.025	.017	.000	.002	.030	.048	.123
HF-CR	.560	.227	.044	.123	.154	.380	1.489

Table 4.3: Total yields of random crossing backgrounds in the signal and HF control regions for the six channels.

4.4 Signal Extraction

As previously mentioned, six channels are used in this analysis, and a global signal strength is extracted from all channels for each model. Depending on the model being considered, an HNL signal may or may not have contributions in a certain channel. For instance, a μ only mixing HNL will have events in the $\mu-\mu\mu$ SR, but not in the $e-\mu\mu$ SR, since the prompt lepton and the displaced leptons have different flavors in the latter. This fact is used to use the SRs in certain channels to help constrain the HF background along with the CRs. Due to the extremely low HF background contribution in the $e-ee$ and $\mu-ee$ SRs, they are combined in the fit to an $x-ee$ SR. Table 4.4 shows the primary use of the SRs and HF-CRs defined in the analysis for each channel considered. This section outlines the uncertainties considered in the analysis and the structure of the statistical analysis used to extract the global signal strength.

4.4.1 Uncertainties

All relevant sources of uncertainties are considered in this analysis which can affect the estimate on the background and signal MC predictions. One standard deviation ($\pm 1\sigma$) variations are considered for all uncertainties, and are propagated to the statistical fit. The standard-

Model	Signal Region					Heavy Flavor Control Region
	$\mu-\mu\mu$	$\mu-\mu e$	$e-\mu\mu$	$e-e\mu$	$x-ee$	All Channels
1SFH(μ)	Signal	Signal	Background	Background	Signal	Background
1SFH(e)	Background	Background	Signal	Signal	Signal	
2QDH(IH)	Signal	Signal	Signal	Signal	Signal	
2QDH(NH)	Signal	Signal	Signal	Signal	Signal	

Table 4.4: A summary of the use of the fit regions defined for the models that are used in the interpretation of the analysis results. Signal represents when a region is mainly used for signal extraction, and Background represents when a region is used for constraining the background from Heavy Flavor decays.

ization of LRT objects in this version of the analysis allowed for a better understanding of the uncertainties on these objects in close coordination with the Combined Performance (CP) groups. The use of MC simulations for the modeling of HF decay background warranted an inclusion of the theoretical modeling uncertainties on them, which have been estimated using the prescription from the Physics Modeling Group. This section describes all the sources of systematic uncertainties considered in this analysis. The net impact of the systematic uncertainties on the analysis is shown in table 5.3. Systematics are found the be a sub-dominant source of uncertainties and the results are limited by the statistical size of the dataset.

Integrated Luminosity

The LUCID2 detector was used to measure the integrated luminosity collected by ATLAS during Run-2 with a central value of 140.1 pb^{-1} and an error of 0.83% [85].

Pileup

While pileup is modeled in simulation, the simulated collision conditions are not identical to those experienced during the data collection. A centrally managed tool is used to re-weight signal and background MC events and take into account the impact of pileup during the various data periods.

Trigger

The efficiency of triggering by a muon or an electron in MC is adjusted to that in data by looking at $Z \to \ell\ell$ events. The efficiency calibrations from the CP groups are used to account for these uncertainties on the primary triggering object, which is the prompt lepton in this analysis.

Lepton Momentum

The momentum of the leptons measured in MC are calibrated for detector conditions using Z and J/ψ decays. The standard CP efficiency calibrations are used for all leptons used in this analysis.

Lepton Reconstruction and Identification

The muons in this analysis are either standard or LRT, all accepted using the `Medium` identification WP. For standard muons, the standard CP efficiency calibrations are applied obtained using $Z(J/\psi)$ decays for muon $p_{\mathrm{T}} > (<)\, 10$ GeV in a tag and probe analysis [55]. For LRT muons, the reconstruction and identification efficiencies are studied separately and are measured indirectly. The efficiency to measure a `Medium` muon can be divided into two factors:

$$\epsilon\,(\text{Medium } \mu) \approx \epsilon\,(\text{Medium } \mu \,|\, \text{ID}) \times \epsilon(\text{ID}), \tag{4.4}$$

Here, $\epsilon\,(\text{LRT ID})$ is roughly estimated using $K_S \to \pi^+\pi^-$ decays [4]. Furthermore, it is postulated that

$$\epsilon\,(\text{Medium LRT } \mu \,|\, \text{LRT ID}) \approx \epsilon\,(\text{Medium Standard } \mu \,|\, \text{Standard ID}), \tag{4.5}$$

since the MS is *far* from the IP, and hence is insensitive to cm level displacements that come from the difference in Standard and LRT ID track reconstruction. The above two inferences allow for the calibration of LRT muon reconstruction and identification efficiencies using the same Z and J/ψ decays, since $\epsilon\,(\text{Medium Standard } \mu\,|\,\text{Standard ID})$ is measured by the standard tag and probe method with the ID track corrections (i.e. contribution from $\epsilon\,(\text{Standard ID})$) disabled in the measurement.

After studying the above terms, an uncertainty of 1(3)% is applied on LRT muons with $p_{\mathrm{T}} > (<)\,5$ GeV.

A similar decomposition also applies to the reconstruction and identification of electrons, where separate reconstruction efficiencies are calculated for standard and LRT electrons. For the identification criteria, separate efficiencies are applied to the prompt electron and the displaced electron, since they use the `Medium` and `VeryLooseNoPix` identification WPs, respectively.

The electron reconstruction efficiency is accounted with a 3% uncertainty for both the standard and LRT objects. The `Medium` electron identification efficiency is measured with a 4% uncertainty, and `VeryLooseNoPix` electrons are identified with a 5% uncertainty.

Lepton Isolation

The CP efficiency calibrations are used on the calibration of the isolation applied on the prompt and displaced muons, while a conservative 2% uncertainty is applied on the electron isolation efficiency. Since the same working point is applied on standard and LRT objects without modifications, the same uncertainties are applied to both of them.

Flavor Tagging

Uncertainties from the measurement of the jet energy scale and resolution were found to have no impact on this analysis. The systematic uncertainties from the calibration of the flavor tagging algorithm are considered using the medium reduction scheme, as recommended by the flavor tagging CP group.

Signal Modeling

In the 2022 analysis, the HNL process was modeled at LO using PYTHIA, and the MC prediction was attached a modeling uncertainty of $3\% \oplus 5\%$, with the former coming from the measurement of the W production cross section and the latter from the analytical calculation of the HNL branching fraction compared to the prediction from MADGRAPH5. It was found that the spin correlations of the HNL decay products were not well modeled by PYTHIA, and a dedicated reweighting procedure was applied to adjust for the same. In this analysis, the predictions are modeled at NLO with the spin correlation well modeled by the MC event generators. Hence, a reduced yet conservative uncertainty of 5% is considered on the signal prediction.

Random Crossing Modeling

The random crossing background estimate reuses ID tracks to create a preliminary track pool, which are later used to reconstruct DVs with uncorrelated objects. The statistical correlation coming from such a multi-count is accounted using a bootstrapping [86] mechanism. Furthermore, alternate track sampling methods account for uncertainties coming from the bias in the sampling strategy of such tracks. A conservative 50% uncertainty is applied on the random crossing background estimate to cover the above effects. Given the very small contribution of this background in all the considered regions, this uncertainty is expected to have a negligible effect.

HF Background Modeling

As described in section 3.2, $t\bar{t}$ and V+jets MC samples are used to model the background from HF decays. Theoretical uncertainties are evaluated for both of these MC simulations following the recommendations of the ATLAS Physics Modeling Group:

- $\mathbf{t\bar{t}}$: Uncertainties are evaluated for the hard scatter generation (or matrix element calculation), parton showering process, QCD scale, modeling of the initial and final state radiations, and the choice of the parton density function.

The matrix element uncertainty is evaluated using a two-point comparison between the nominal POWHEG sample showered with PYTHIA 8 and a sample generated with an identical configuration except the `pThard` parameter of POWHEG set to 1 instead of the nominal value of 0. The parton shower uncertainty is evaluated using a two-point comparison between the nominal POWHEG sample showered with PYTHIA 8 and HERWIG,7 showering program.

QCD scale uncertainty is evaluated using the five-point envelope of the variations of renormalization and factorization scales by factors of 2.0 or 0.5 (removing the cases when both are 0.5, both are 2, and one is 2.0 and the other is 0.5). The initial and final state radiation uncertainties are calculated using the $\pm 1\sigma$ weights from PYTHIA 8. Lastly, the PDF uncertainty is evaluated using the NNPDF3.0NLO variations in the POWHEG+PYTHIA 8 sample as defined in the `PDF4LHC` recommendations [87].

- **V+jets**: Uncertainties are evaluated for the QCD scale and parton density function choice.

  QCD scale uncertainty is evaluated using the seven-point envelope of the variations of renormalization and factorization scales by factors of 2.0 or 0.5 (removing the cases when both are 0.5 and both are 2). The PDF uncertainty is evaluated in the same manner as for the $t\bar{t}$ samples following the `PDF4LHC` recommendations.

Table 4.5 summarizes the theory uncertainties used for the background modeling and the method used to calculate them.

HF Template Building

As described previously, a dedicated uncertainty is assigned on the template building process of the HF background. It is equal to the Monte Carlo Statistics uncertainty in the HF-CR, since no smoothing or shape building process is applied in those regions. In the SR, a two-

Process	Uncertainty	Variations	Method
$t\bar{t}$	Matrix Element	POWHEG+PYTHIA 8 `pThard`:0, POWHEG+PYTHIA 8 `pThard`:1	Two-point variation
	Parton Shower	POWHEG+PYTHIA 8, POWHEG+HERWIG 7	Two-point variation
	Initial State Radiation	Var3cUp, Var3cDown	up down variation
	Final State Radiation	muRfac=2.0, muRfac=0.5	up down variation
	QCD Scale	μ_R, $\mu_F \in \{.5, 1, 2\}$	5-point envelope
	PDF	NNPDF30_nlo_as_0118	$\sqrt{\frac{1}{N}\sum_i (X_i - X_0)^2}$
V+jets	QCD Scale	μ_R, $\mu_F \in \{.5, 1, 2\}$	7-point envelope
	PDF	NNPDF30_nlo_as_0118	$\sqrt{\frac{1}{N}\sum_i (X_i - X_0)^2}$

Table 4.5: Summary of the theoretical systematic uncertainties considered on the background Monte Carlo simulations and their method of evaluation.

point uncertainty comparing the ShR shape from a single channel and the ShR shape from the combination of all channels is applied. Since the shapes are normalized to the MC event yield of individual channels, this is a shape-only effect.

Monte Carlo Statistics

Lastly, the MC statistics uncertainty is considered on the signal prediction as well coming from the limited statistics considered in the production of the signal MC samples.

4.4.2 Simultaneous fit

A global likelihood minimizing fit is performed to extract a signal strength for each simulated signal point. A likelihood function $\mathcal{L}$ is used to define a global probability density function of this measurement. The extended likelihood $\mathcal{L}$ consists of a product of likelihood functions ($\mathcal{P}$s) of the SR yield and CR yield. The function is shown without a channel indexing for simplicity, but a global product runs over all five SRs and six CRs considered in the analysis:

$$\mathcal{L} = \mathcal{P}(N^{\text{SR}}_{\text{data}}|\mu \cdot N^{\text{SR}}_{\text{signal}} + \mu_{\text{HF}} \cdot N^{\text{SR}}_{\text{HF}} + N^{\text{SR}}_{\text{RC}}) \cdot \mathcal{P}(N^{\text{CR}}_{\text{data}}|\mu_{\text{HF}} \cdot N^{\text{CR}}_{\text{HF}} + N^{\text{CR}}_{\text{RC}}) \cdot \prod_{i=1}^{n_{\theta_i}} \mathcal{N}(\theta_i|\tilde{\theta}_i), \quad (4.6)$$

where $N_{\text{data}}^{\text{SR}}$ and $N_{\text{data}}^{\text{CR}}$ are the numbers of events observed in data in the signal region and the HF control region, respectively, and $N_{\text{signal}}^{\text{SR}}$ is the expected signal yield in the SR. The parameter μ is the normalization of the signal, and μ_{HF} that of the HF background. A signal strength of $\mu = 0$ corresponds to a background only hypothesis while $\mu = 1$ corresponds the nominal signal plus background (S+B) hypothesis. $N_{\text{HF}}^{\text{SR}}$ and $N_{\text{HF}}^{\text{CR}}$ are the expected numbers of HF background events in the SR and CR, respectively. As previously mentioned, the RC background is taken directly from the data-driven prediction and is represented by $N_{\text{RC}}^{\text{SR}}$ and $N_{\text{RC}}^{\text{CR}}$ in the SR and CR, respectively. The term $\mathcal{N}$ is a normal probability distribution function of the systematic effect of the type i, represented by the nuisance parameter θ_i, around its nominal value $\tilde{\theta}_i$. The parameter n_{θ_i} represents the number of nuisance parameters modeling the systematic uncertainties considered in this analysis.

The likelihood function $\mathcal{P}$ in the SR and CR are defined as Poisson probability density functions [88] built from binned histograms (also referred to as templates) of m_{HNL} and m_{DV}, respectively, using MC events for signal and HF background and a data driven template for the RC background. Bins of 500 MeV from 0 to 20 GeV in m_{HNL} are used in the SR determined by the m_{HNL} reconstruction resolution in signal. Similarly, bins of 500 MeV from 0 to 6 GeV in m_{DV} are used in the CR, since no HF background is observed beyond 6 GeV. The signal and background templates are also functions of the nuisance parameters [88] such that systematic uncertainties can affect the m_{HNL} or m_{DV} shapes for the signal process and the HF background whose yields are extracted from the fit and the normalization for the RC background whose yield is fixed to the prediction. This methodology accounts for correlations of systematic uncertainties between signal and background estimates. All systematics considered in the analysis are fully correlated between regions and whichever samples they apply to, expect for the systematics from template building and MC statistics, which are decorrelated between each bin and region.

The likelihood is maximized to determine the best fit value of the global signal strength μ. More specifically, the profile likelihood is used instead of the likelihood to correctly account for the effect of the systematics near the minimization point, and is given by:

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{\theta}}_\mu)}{\mathcal{L}(\hat{\mu}, \hat{\theta}_\mu)}, \tag{4.7}$$

where $\hat{\mu}$ and $\hat{\theta}_\mu$ are estimators for the signal strength and nuisance parameters that maximize the likelihood globally, and $\hat{\hat{\theta}}_\mu$ is a conditional likelihood maximizing estimator for a specific μ.

The test statistic is built using the profile likelihood ratio which is minimized to determine the best fit value:

$$q_\mu = -2 \ln \lambda(\mu). \tag{4.8}$$

The probability of the background fluctuating such that it is equal to larger than the number of data events is given by the p-value p_0:

$$p_0 = \int_{q_{\text{obs}}}^{\infty} f\left(q_0 \mid 0, \hat{\theta}_0\right) \mathrm{d}q_0, \tag{4.9}$$

where the integrand is the probability distribution function in the background-only hypothesis q_0 and q_{obs} is the observed test statistic. Similarly, a p-value for the S+B hypothesis p_μ is defined as:

$$p_\mu = \int_{q_{\text{obs}}}^{\infty} f\left(q_\mu \mid \mu, \hat{\theta}_\mu\right) \mathrm{d}q_\mu. \tag{4.10}$$

The CL$_s$ [89] method is used to exclude signal hypothesis since it has been shown to perform better than a standard p-value based exclusion in analysis of statistically limit datasets. The test statistic is defined as:

$$\text{CL}_s = \frac{p_\mu}{p_0}. \tag{4.11}$$

The 95% confidence level (CL) upper limit on the signal strength can be found by varying μ until $\text{CL}_s(\mu_{\text{limit}}) = 0.05$. Signal models are excluded if $\mu_{\text{limit}} < 1$.

Figure 4.18 shows the m_{DV} distribution of measured data, random crossing background, and expected HF background in the six CRs. Good agreement between the data and predictions is observed across all regions within errors in data counting (shown by black lines around data points) and systematic uncertainties (shown by hashed bands around the HF background estimate), with some discrepancies expected to be corrected by the fit.

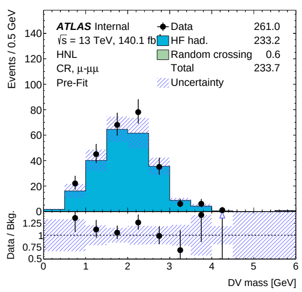

(a) HF-CR $\mu - \mu\mu$

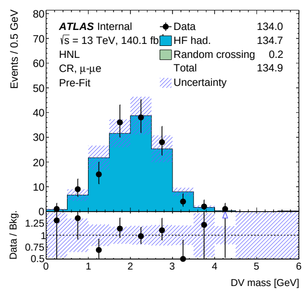

(b) HF-CR $\mu - \mu e$

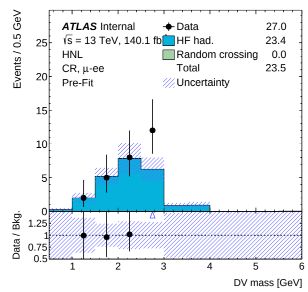

(c) HF-CR $\mu - ee$

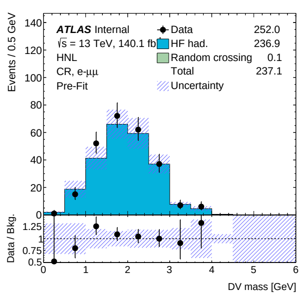

(d) HF-CR $e - \mu\mu$

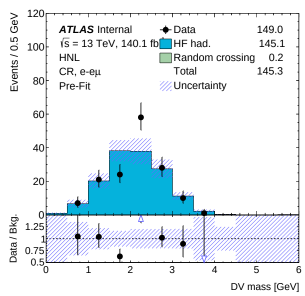

(e) HF-CR $e - e\mu$

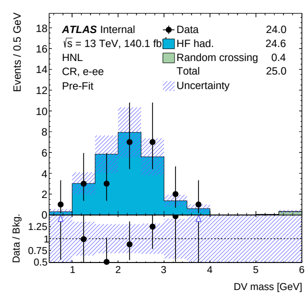

(f) HF-CR $e - ee$

Figure 4.18: Data and pre-fit background m_{DV} distributions in the six control regions. Inclusive yields shown in the legend are rounded to the first decimal point.

Similarly, fig. 4.19 shows the m_{HNL} distributions of measured data and expected backgrounds in the five SRs. A good compatibility is also seen between the background predictions and the observed data, with the exception of a few outlying events, an analysis of which is presented in section 5.1.

4.5 Background-Only Fits

Background-only fits are performed to ascertain the stability of the fit structure and test the impact of the considered sources of systematic errors. Two kinds of background-only fits are

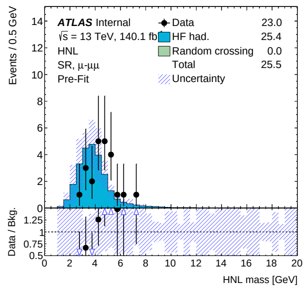

(a) SR $\mu - \mu\mu$

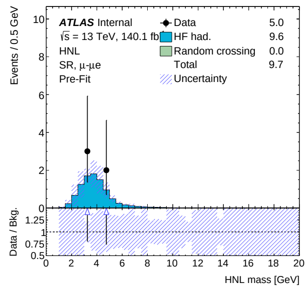

(b) SR $\mu - \mu e$

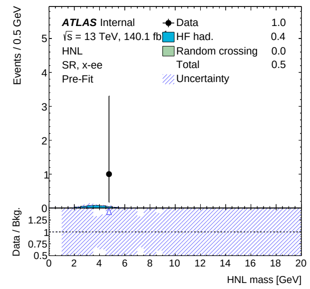

(c) SR $x - ee$

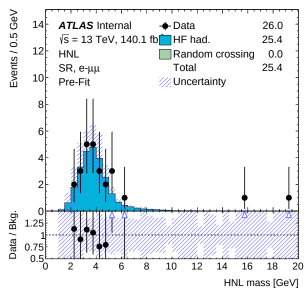

(d) SR $e - \mu\mu$

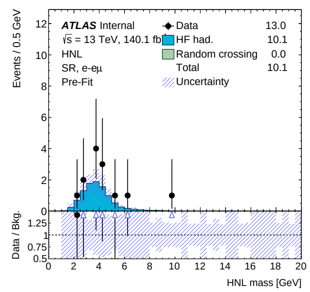

(e) SR $e - e\mu$

Figure 4.19: Data and pre-fit background $m_{\rm HNL}$ distributions in the five signal regions. Inclusive yields shown in the legend are rounded to the first decimal point.

performed:

- Asimov fit: An Asimov (pseudo-)dataset [90], built using the expected number of events, is used to test the basic structure of the simultaneous fit, i.e. its ability to reduce the impact of nuisance parameters on the background prediction. Figure 4.20 shows the impact of the fit in the $\mu - \mu\mu$ control region. After the fit, an HF normalization factor $\mu_{\rm HF} = \mathbf{1.00^{.11}_{-.10}}$ is observed. An error of the order of 10% on the measurement shows a reduction in errors from the fit in comparison to pre-fit, and a powerful ability to estimate this background. This reduction in errors improves the measurement sensitivity with respect to a simple cut and count analysis.

  Figure 4.21a shows the constraints of the nuisance parameters considered (except for template building systematics) in the analysis. None of the experimental systematics are significantly constrained compared to their pre-fit 1σ values. On the other hand, the background modeling systematics (PS-parton shower, ME-matrix element, FSR-

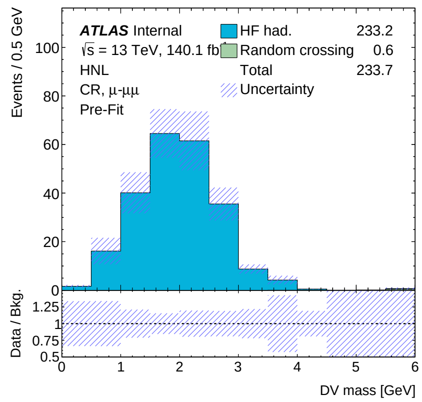

(a) Pre-fit

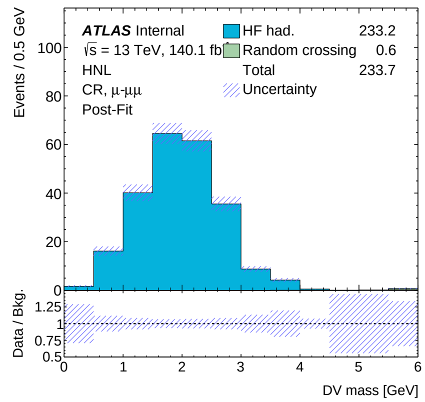

(b) Post-fit

Figure 4.20: Pre- and post-fit background distributions in the $\mu - \mu\mu$ control region. The post-fit distribution shows a significant reduction in errors compared to the pre-fit inputs.

final state radiation) are seen to be constrained. While is expected given their relatively large pre-fit uncertainties, especially in the signal regions, the constraints imply that the variation model recommended by the collaboration may not be sufficiently granular for estimating the HF background using these MC simulations. Considering the relatively small impact of these systematics on the overall measurement, the constraints are deemed acceptable for the purpose of estimating this background in this analysis and the net effect is a reduction in overall errors.

- Data fit: A background-only fit is performed using measured ATLAS data in both the signal and control regions. Such a fit tests the compatibility of the of the Asimov and data fit, in the capacity of treatment of systematics as well as that of the background normalization factor, and gives an indirect sign of any excesses beyond expectation in the region where the background is concentrated. An HF normalization factor $\mu_{\text{HF}} = \mathbf{1.04^{.12}_{-.10}}$ is observed, which is highly compatible with what is measured by the Asimov background-only fit with comparable errors. Figure 4.21b shows the pulls and constraints of the nuisance parameters in the data fit. The constraints are very similar to what is seen in the Asimov fit. Some pulls, which represent the compatibility of θ measurements between $\mu_{\text{HF}} = 0$ and the measured μ_{HF}, are observed. Furthermore,

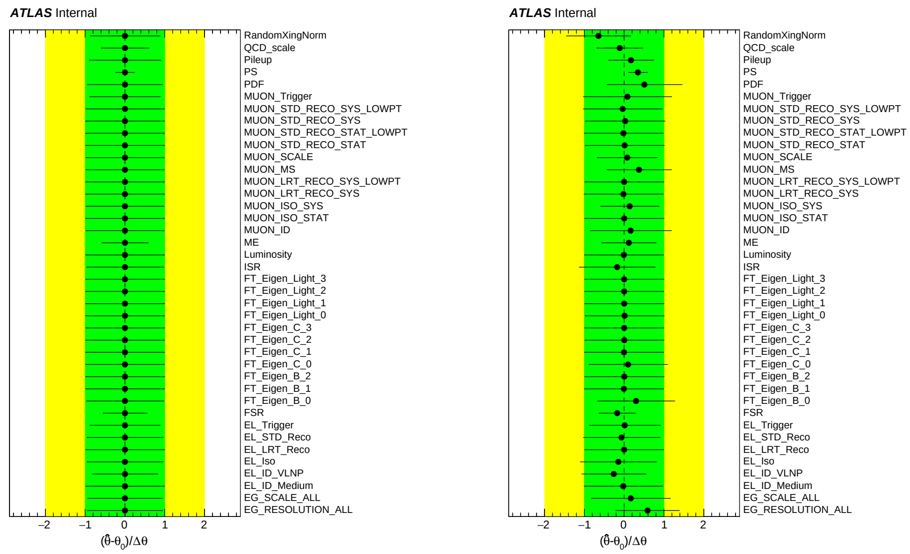

(a) Asimov fit

(b) Data fit

Figure 4.21: Pulls and constraints of the systematic errors considered in the fit (in no particular order).

the template building systematic in bins 32 ($m_{\mathrm{HNL}} = 15.5 - 16$ GeV) and 39 ($m_{\mathrm{HNL}} = 19 - 19.5$ GeV) of SR $e - \mu\mu$ is pulled by $+.6\sigma$. This is expected because of the 2 excess data event observed in those bins which may under-predicted by the MC simulations, or the events could just be random fluctuations given the limited size of the dataset.

Overall, the background-only fits using the Asimov and the real dataset show that the background prediction is robust in the regions where the MC simulations predict events. The 1 extra event seen in SR $e - e\mu$ and and 2 extra events seen in SR $e - \mu\mu$ lie in the tails of the m_{HNL} spectrum where the statistics in the MC simulations are not sufficient for a highly granular prediction.

Chapter 5

RESULTS

This chapter presents the results of this analysis. Section 5.1 presents a detailed analysis of the three outlying events seen in the signal regions of the channels $e - \mu\mu$ and $e - e\mu$. The outcome of the search for HNLs is presented in section 5.2 in the form of excluded phase spaces and comparisons to a previous search.

5.1 Analysis of Outlying Events

As shown in fig. 4.19, all but three data events are found to be compatible (within statistical and systematic coverage) with the MC prediction. Two of these three events are found in the $e - \mu\mu$ SR while the third is found in the $e - e\mu$ SR. Table 5.1 tabulates the important features and kinematics of the these events.

Outlier Event	1	2	3
Channel	$e - \mu\mu$	$e - \mu\mu$	$e - e\mu$
m_{HNL} [GeV]	19.1	15.8	9.6
m_{DV} [GeV]	2.70	1.9	2.8
$\mathcal{S}$	247.8	128.5	124.4
Prompt lep. p_{T} [GeV]	28.5	44.8	45.4
$\Delta\eta_{\ell\ell}$	0.11	0.10	0.09
$\Delta\phi_{\ell\ell}$ [rad]	0.22	0.19	0.16
$\Delta R_{\ell\ell}$	0.24	0.22	0.19
Cosmic Separation	5.6	4.2	3.7
$\theta_{r_{\text{DV}} p_{\text{DV}}}$ [rad]	0.06	0.13	0.08
Data Taking Year	2015	2018	2018
Run Number	283429	356250	350184
Event Number	1139948317	1118382313	2746409499

Table 5.1: Relevant kinematics and features of the three outlying events found in the analysis signal regions.

The m_{DV} values of the events lie within 1-3 GeV, typically dominated by the HF background. The opening angle of the DV ($\Delta R_{\ell\ell}$) and its azimuthal and polar projections are relatively small, indicating that the DVs are most likely from genuine decays rather than random crossings. A very large value of the cosmic separation variable also rules out the possibility that these events are from cosmic muons. Lastly, the angle between the position of the DV and its momentum ($\theta_{r_{\rm DV}p_{\rm DV}}$) is very small, indicating that there are no losses in the visible DV momentum in the form of neutrinos. The above facts suggest that these events are not predicted by the background estimate due to an underpopulated or a missing Monte Carlo sample, perhaps QCD multi-jet, which leads to the same final state and contaminates the measured phase space. An HNL signal would show signatures in multiple channels in similar mass windows, and hence would give a statistically significant deviation from the SM expectation.

A statistical discovery search is performed on the full set of regions to go beyond the 'by-eye' evaluation. The p-values are calculated for the four signal models and for the $c\tau_{\rm HNL}$ =10 mm and 100 mm points. One thousand pseudo-experiments are performed for the signal and signal-plus-background hypotheses. The test-statistic used for this study is the *uncapped* variant of q_μ:

$$q_\mu^{uncap} = \begin{cases} -2\ln\frac{\mathcal{L}(\mu=0,\hat{\hat{\theta}}_{\mu=0})}{\mathcal{L}(\hat{\mu},\hat{\theta}_\mu)} & \text{for } \hat{\mu} \geq 0 \\ +2\ln\frac{\mathcal{L}(\mu=0,\hat{\hat{\theta}}_{\mu=0})}{\mathcal{L}(\hat{\mu},\hat{\theta}_\mu)} & \text{for } \hat{\mu} < 0 \end{cases}, \tag{5.1}$$

which takes positive values for $\hat{\mu} > 0$ and negative values for $\hat{\mu} < 0$. The signal significance Z is hence computed as:

$$Z = \begin{cases} +\sqrt{+q} & \text{for } \hat{\mu} \geq 0 \\ -\sqrt{-q} & \text{for } \hat{\mu} < 0 \end{cases}. \tag{5.2}$$

Hence, such a statistic allows for negative signal significances as well, arising from negative post-fit signal yields. A negative signal fluctuation could arise out of fractional background predictions with zero data events in a bin, or as artefacts from bins with zero background as well as data. The results from the study are shown in fig. 5.1 for signal with $c\tau_{\rm HNL} = 10$ and 100 mm. Negative significances are shifted to zero (p-value = 0.5) for illustration.

The discovery significances across the mass spectrum are seen to be compatible for the

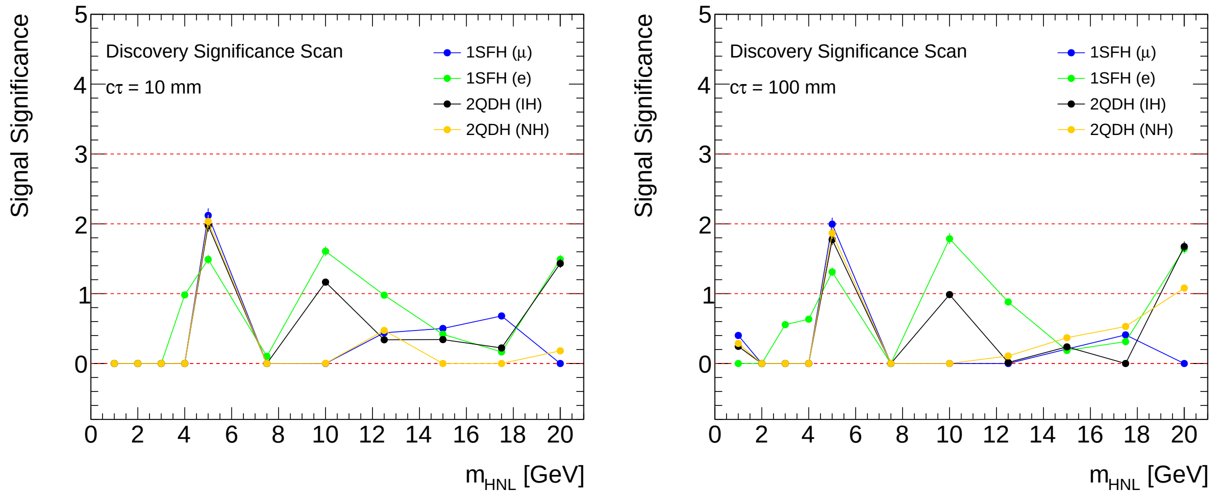

(a) $c\tau_{\mathrm{HNL}}$= 10 mm

(b) $c\tau_{\mathrm{HNL}}$= 100 mm

Figure 5.1: Discovery significance as a function of mass for two representative HNL lifetimes for the four models considered. Negative significance values are shifted to zero.

10 mm and 100 mm samples accounting for shape differences in the signal template from acceptance and efficiency effects of analysis cuts. There are three main features of the observed data events and the predicted background yields (see fig. 4.19) in the signal region which affect and explain the observed discovery significance:

1. Under-fluctiation of background prediction: No data is observed in the SR of all channels up to m_{HNL} = 2.5 GeV and in $m_{\mathrm{HNL}} = 5 - 8$ GeV in the $\mu - \mu e$ SR. The non-negligible background predictions in these regions cause a $\hat{\mu} < 0$ to be possible, and hence negative significances are observed in models that leave signatures in these mass ranges in the relevant channels.

2. Outliers described in table 5.1: The m_{HNL} = 19.1 GeV event in the $e - \mu\mu$ channel creates a positive significance at the level of 1.5 standard deviations in the 1SFH(e) and 2QDH(IH) model, since the 1SFH(μ) model does not have signal events in the $e - \mu\mu$ channel and the contribution from the 2QDH(NH) is also low, since it has a much higher coupling to muons than electrons. The m_{HNL} = 16 GeV event does not cause a significant excess, since the simulated points are at 15 and 17.5 GeV, which

only have their mass tails at the outlying point. The last outlier event at 9.6 GeV in the $e - e\mu$ channel also creates a positive significance of about 1.5 standard deviations in the 1SFH(e) and 2QDH(IH) models, while the other two models that are largely coupled to muons are not sensitive to this event.

3. Excess around $m_{\mathrm{HNL}} = 5$ GeV: In all channels at varying degrees, a slight pre-fit underestimation of data by the background estimate around $m_{\mathrm{HNL}} = 4 - 6$ GeV is observed. This excess of data causes a signal significance observation of 2 standard deviations for all models except 1SFH(e), which is at the level of 1.5 standard deviations. A possible reason for such an effect is the shape building process of the HF background in the SR, where all channels are combined to build the template shape, with the underlying educated assumption that the HF background has a similar signature irrespective of the lepton flavor.

The discovery significance study shows **no deviation beyond statistical fluctuations** in this measurement, even though there are isolated excesses which could have been significant. This shows the strength of the analysis method where the signal is extracted using multiple regions at once using the m_{HNL} shape, where both of these techniques (simultaneous fit and the m_{HNL} shape) reduce the influence of such fluctuations on the overall measurement.

Channel	Signal Region			Control Region		
	Hadrons	Crossings	Obs.	Hadrons	Crossings	Obs.
$\mu - \mu\mu$	23.1 ± 2.6	0.023 ± 0.012	23	254.4 ± 13.4	0.16 ± 0.08	261
$\mu - \mu e$	9.2 ± 0.8	0.016 ± 0.008	5	135.3 ± 7.9	0.096 ± 0.052	134
$\mu - ee$	-	-	-	23.7 ± 2.5	0.015 ± 0.008	27
$e - \mu\mu$	25.9 ± 1.7	0.0014 ± 0.0007	26	248.8 ± 12.1	0.051 ± 0.028	252
$e - e\mu$	9.6 ± 0.9	0.027 ± 0.015	13	153.3 ± 8.7	0.057 ± 0.030	149
$e - ee$	-	-	-	22.4 ± 2.8	0.094 ± 0.050	24
$x - ee$	0.22 ± 0.14	0.043 ± 0.023	1	-	-	-

Table 5.2: Post-fit yields of heavy flavor hadron decay background (Hadrons), random crossing background (Crossings), and data (Obs.) in the signal and control regions using the $m_{\mathrm{HNL}} = 5$ GeV, $c\tau_{\mathrm{HNL}} = 10$ mm 2QDH(IH) model as a benchmark.

Table 5.2 shows the post-fit yields for the observed data and background processes obtained using the $m_{\mathrm{HNL}} = 5$ GeV, $c\tau_{\mathrm{HNL}} = 10$ mm samples in the 2QDH(IH) model as a benchmark. The simultaneous fit gives a signal strength of $\hat{\mu} = 0.041 \pm 0.024$ and a HF background normalization of $\hat{\mu_b} = 1.16 \pm 0.11$. This model is easily excluded, as is shown in the next section.

Table 5.3 shows the percent impact of the sources of uncertainties considered in this analysis obtained using the 2QDH(IH) model for HNL masses and lifetimes at the edge of the exclusion contour (see section 5.2). As expected, the random crossing estimate becomes relevant only at larger m_{HNL}, since the HF background dominates at relatively lower masses. At very low masses where there is no contribution from the HF background, the signal modeling and lepton measurement systematics dominate. Overall, however, the impact of systematic uncertainties is significantly smaller than the statistical uncertainty.

Uncertainty Source	Impact (%) on $\hat{\mu}$ using 2QDH(IH) (m_{HNL}, $c\tau_{\mathrm{HNL}}$)			
	1 GeV, 1000 mm	3 GeV, 1 mm	5 GeV, 10 mm	12.5 GeV, 10 mm
Leptons	2.7	2.0	4.5	2.7
Flavor Tagging	< 0.1	1.3	1.5	0.4
Pileup	1.1	1.1	3.3	< 0.1
Luminosity	0.3	0.3	0.6	0.3
Signal Modeling	1.7	1.8	3.4	1.8
HF Bkg. Modeling	0.2	9.1	9.6	1.1
RC Bkg. Modeling	< 0.1	< 0.1	< 0.1	24.2
HF Bkg. Template	0.4	2.9	3.9	1.1
Signal MC Stats.	12.2	3.2	3.4	3.1
HF Normalization	1.2	2.4	1.3	1.3
Total Syst.	12.7	11.0	12.6	4.8
Data Stats.	94.7	112.5	33.0	175.5
Total	95.5	113.0	35.3	175.6

Table 5.3: Percentage impact of uncertainties on the measured signal strength $\hat{\mu}$ for the 2QDH(IH) HNL model for representative mass and proper lifetimes at the edges and in the bulk of the sensitivity.

5.2 Exclusion Limits

Since no deviations beyond statistical fluctuations are found, 95% CL upper limits of the signal strength are obtained for the four benchmark models, and the exclusion contours are shown in fig. 5.2. The limits are obtained using the "Asymptotic" method [90] which performs an analytical calculation using a χ^2 approximation of q_μ. A comparison of using the Aymptotic method against pseudo-experiments is shown in chapter C of the appendix. The figures show the observed and expected exclusion contours as a function of the HNL mass and the square of mixing to SM neutrinos ($|U|^2$), and the same is compared to the expected exclusion from the 2022 analysis. All the mass and couplings within the contour are excluded from possible HNL models.

The shape the contour is well explained by the relationship between $|U|^2$, the HNL proper lifetime, and the HNL mass, $\tau_{\mathrm{HNL}} \propto \frac{1}{m_{\mathrm{HNL}}^5 |U|^2}$. The analysis shows sensitivity to models with $c\tau_{\mathrm{HNL}} = 1 - 1000$ mm, corresponding to mixing angles as low as 10^{-6} up to $m_{\mathrm{HNL}} = 5$ GeV. HNLs with larger lifetimes do not have enough decays in the ID volume for the analysis to be sensitive to them. HNLs with shorter lifetime fail the $r_{\mathrm{DV}} > 4$ mm requirement since they are not displaced enough to be captured by the secondary vertex selection. As the mass increases, the HNL production cross-section reduces, decreasing the overall sensitivity to these HNLs. The 1SFH(μ) contour shows an impressive level of exclusion where a 16 GeV HNL is excluded for mixing angles as low as 10^{-7}, a region untouched by other measurements. The contour is open in the low mass end since there is no estimate of the level of exclusion for $m_{\mathrm{HNL}} < 1$ GeV. To provide measurable (but not unphysically) large couplings, HNLs with smaller (< 1 GeV) masses require very large proper lifetimes to which the analysis has no acceptance. The $m_{\mathrm{HNL}} = 1$ GeV, $c\tau_{\mathrm{HNL}} = 1000$ mm HNLs are excluded for all models, a lifetime not excluded for any mass in the 2022 analysis. The expected and observed exclusions are compatible with each other. However, the observed exclusion has a slight dip in the $m_{\mathrm{HNL}} = 6 - 8$ GeV region. This is precisely because of the excess/peak shift effect around 5 GeV that was described in section 5.1, where a larger amount of data is observed than the background prediction. Nevertheless, no excess beyond two-standard deviation level is measured, as was also verified by the discovery significance scans.

Figure 5.3 shows a simplified comparison of the expected exclusion from this result

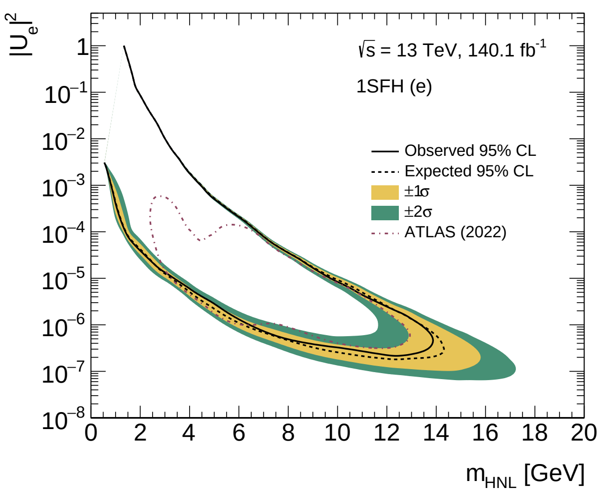

(a) 1SFH(e)

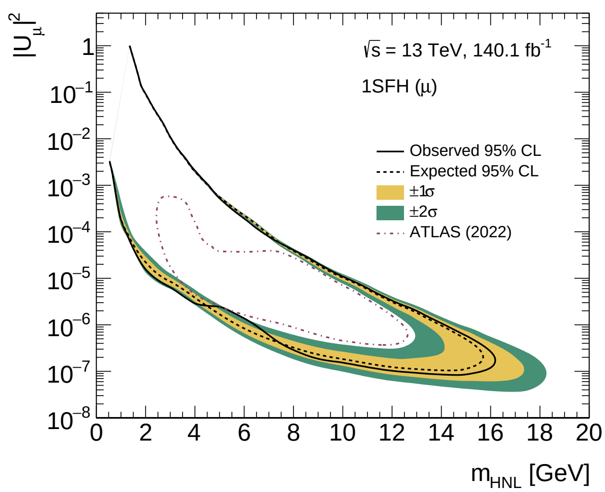

(b) 1SFH(μ)

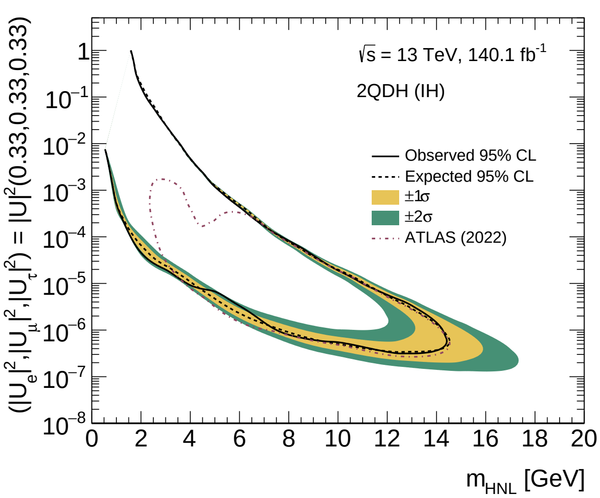

(c) 2QDH(IH)

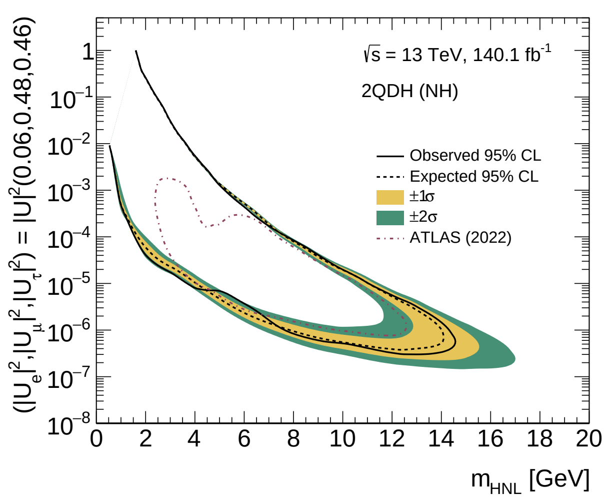

(d) 2QDH(NH)

Figure 5.2: The observed (solid line) and expected (dashed line) 95% exclusion contours for the four benchmark models as a function of the HNL mass and $|U|^2$. The green and yellow bands shows the range of limits within 1 and 2 standard deviations, respectively. The violet dashed line shows the expected exclusion contour of the 2022 analysis.

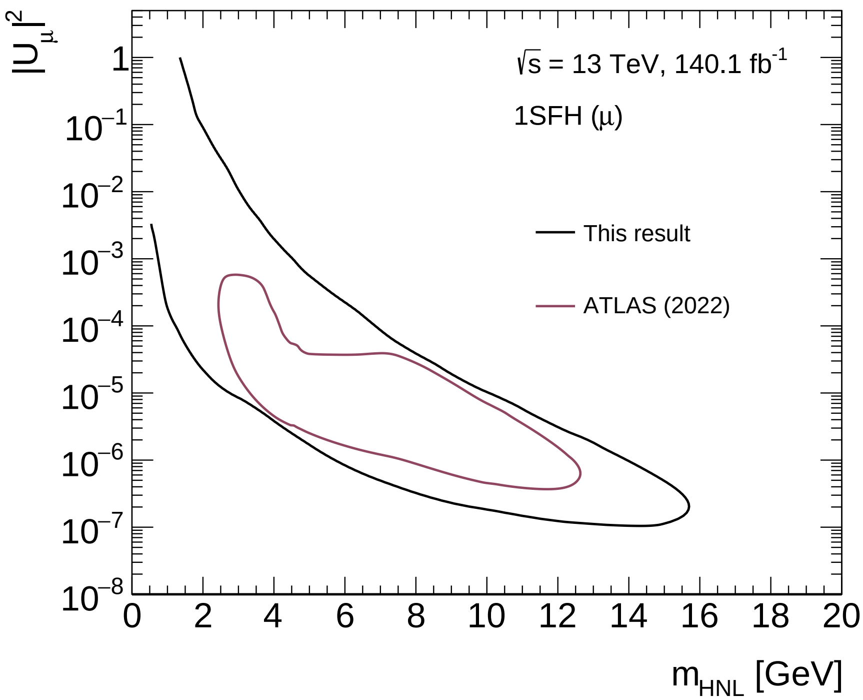

Figure 5.3: The expected 95% exclusion contours for the 1SFH(μ) HNL model as a function of the HNL mass and $|U|^2$ for this result (black) and the 2022 analysis (violet).

compared to the 2022 analysis. The improved reconstruction and analysis methods used in this analysis with respect to the 2022 analysis are apparent in the comparison. The first visible improvement is an extension of the sensitivity reach to the 1-3 GeV range. The 2022 analysis was forced to apply stringent cuts to reject the HF background which also added the "kink" in their exclusion shape. The removal of this cut and the robust HF background estimation in my analysis allows for the significant improvement to sensitivity in the low mass region. For some models, the high mass exclusion extends up to 16 GeV, to be compared to 14 GeV in the 2022 analysis, thanks to the improvements to VSI (especially a bug fix which resulted in significantly lower random crossings), the improved LRT, and the use of m_{HNL} shape as a discriminant. The most significant difference between the two analyses is in the 1SFH(μ) model, where the exclusion contour is visibly expanded in all directions. The lowest mixing angle excluded by the 2022 analysis using the 1SFH(μ) model was at 12 GeV. This exclusion has been improved by an order of magnitude.

Chapter 6

CONCLUSION

The continued development and improvement of dedicated reconstruction techniques to search for LLPs at the ATLAS experiment makes *now* a very unique time to explore unconventional signatures left by decaying LLPs. The analysis presented in this dissertation searches for Heavy Neutral Leptons decaying into a pair of charged leptons reconstructed as a secondary vertex in the Inner Detector volume. It is sensitive to a wide range of HNL masses, 1 GeV $< m_{\mathrm{HNL}} <$ 16 GeV, and to HNL-ν mixing angles as low as 10^{-7}. The results are a direct improvement on the 2022 ATLAS search [60] that probed the same final state and dataset but used older reconstruction and analysis techniques. The analysis presented is the first in ATLAS using a MC-driven shape estimate for the background estimation in an LLP search, setting a precedent for future searches. No statistically significant deviation from the Standard Model prediction was observed beyond the statistical and systematic precision of the search. The sensitivity reach in the form of an exclusion contour is significantly expanded from all directions, with the results being statistically limited. The exclusion limits obtained in this search are the best-ever obtained using the ATLAS experiment, and provide the world leading exclusions on HNL-ν_μ mixing in the 10-16 GeV HNL mass range. A paper documenting these results is currently in preparation.

Outlook

This analysis is largely statistically limited, meaning that additional data would significantly improve the sensitivity reach. The HL-LHC program provides the ultimate opportunity to collect such a dataset and extend the sensitivity of the analysis. Apart from larger statistics, the sensitivity to lower HNL couplings can also be extended using complementary analysis techniques and final states, a direction that can be explored using the new Run-3 ATLAS dataset. Specifically for lower masses, the search for HNLs decaying semi-leptonically into a lepton and a pion could help add sensitivity from a completely non-overlapping channel.

This extension is natural given the two-track nature of the DV which has already been studied in great detail in this analysis. The DV reconstruction in this analysis uses ID tracks, which puts a hard cut-off at 300 mm on the radial displacement of the DV. Being limited to this volume restricts the HNL lifetime reach of the analysis, since the event yield from weakly coupled HNLs became excessively small in this volume with the 140.1 fb^{-1} dataset. Hence, a more significant extension to the sensitivity can come from increasing the DV acceptance beyond the innermost silicon strips layer. Specifically, vertex reconstruction using clusters in the calorimeter or MS tracks extends the reach to $c\tau_{\text{HNL}} = 1 - 10$ m. Such a reconstruction, however, is more challenging since the calorimeter and the MS don't have the same granularity as the ID in the ϕ-direction, making it difficult to separate genuine displaced decays from random crossings. Lastly, the limited triggering scheme used in this analysis added an implicit requirement on the p_{T} thresholds of the final state leptons restricting the signal acceptance from lower HNL masses. This scheme could possibly be expanded in the future making use of multi-lepton triggers if the soft QCD background (which dominates in the lower lepton p_{T} phase space) can be estimated well.

REFERENCES

[1] D. Curtin et al. "Long-lived particles at the energy frontier: the MATHUSLA physics case". In: *Reports on Progress in Physics* 82.11 (Oct. 2019), p. 116201. DOI: 10.1088/1361-6633/ab28d6. URL: https://dx.doi.org/10.1088/1361-6633/ab28d6.

[2] J. Alimena et al. "Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider". In: *Journal of Physics G: Nuclear and Particle Physics* 47.9 (Sept. 2020), p. 090501. DOI: 10.1088/1361-6471/ab4574. URL: https://dx.doi.org/10.1088/1361-6471/ab4574.

[3] ATLAS Collaboration. *Software Performance of the ATLAS Track Reconstruction for LHC Run 3*. 2023. arXiv: 2308.09471 [hep-ex].

[4] ATLAS Collaboration. "Performance of the reconstruction of large impact parameter tracks in the inner detector of ATLAS". In: *Eur. Phys. J. C* 83 (2023), p. 1081. DOI: 10.1140/epjc/s10052-023-12024-6. arXiv: 2304.12867.

[5] F. Halzen and A. Martin. *Quarks & Leptons: An introductory course in modern particle physics*. New York, USA: John Wiley & Sons, 1984.

[6] S. N. Gninenko, D. S. Gorbunov, and M. E. Shaposhnikov. "Search for GeV-Scale Sterile Neutrinos Responsible for Active Neutrino Oscillations and Baryon Asymmetry of the Universe". In: *Advances in High Energy Physics* 2012 (2012), pp. 1–17. ISSN: 1687-7365. DOI: 10.1155/2012/718259. URL: http://dx.doi.org/10.1155/2012/718259.

[7] LHCb Collaboration. "Observation of the Resonant Character of the $Z(4430)^-$ State". In: *Phys. Rev. Lett.* 112 (22 June 2014), p. 222002. DOI: 10.1103/PhysRevLett.112.222002. URL: https://link.aps.org/doi/10.1103/PhysRevLett.112.222002.

[8] LHCb Collaboration. "Observation of a $J/\psi\Lambda$ Resonance Consistent with a Strange Pentaquark Candidate in $B^- \to J/\psi\Lambda\overline{p}$ Decays". In: *Phys. Rev. Lett.* 131 (3 July 2023), p. 031901. DOI: 10.1103/PhysRevLett.131.031901. URL: https://link.aps.org/doi/10.1103/PhysRevLett.131.031901.

[9] P. W. Higgs. "Broken Symmetries and the Masses of Gauge Bosons". In: *Phys. Rev. Lett.* 13 (16 Oct. 1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.508.

[10] F. Englert and R. Brout. "Broken Symmetry and the Mass of Gauge Vector Mesons". In: *Phys. Rev. Lett.* 13 (9 Aug. 1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.321.

[11] BASE Collaboration. "A 16-parts-per-trillion measurement of the antiproton-to-proton charge–mass ratio". In: *Nature* 601.7891 (Jan. 2022), pp. 53–57. ISSN: 1476-4687. DOI: 10.1038/s41586-021-04203-w. URL: http://dx.doi.org/10.1038/s41586-021-04203-w.

[12] D. Hanneke, S. Fogwell Hoogerheide, and G. Gabrielse. "Cavity control of a single-electron quantum cyclotron: Measuring the electron magnetic moment". In: *Phys. Rev. A* 83 (5 May 2011), p. 052122. DOI: 10.1103/PhysRevA.83.052122. URL: https://link.aps.org/doi/10.1103/PhysRevA.83.052122.

[13] ATLAS Collaboration. *Measurement of the W-boson mass and width with the ATLAS detector using proton-proton collisions at $\sqrt{s} = 7$ TeV*. 2024. arXiv: 2403.15085 [hep-ex].

[14] ATLAS Collaboration. "A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery". In: *Nature* 607.7917 (July 2022), pp. 52–59. ISSN: 1476-4687. DOI: 10.1038/s41586-022-04893-w. URL: http://dx.doi.org/10.1038/s41586-022-04893-w.

[15] CMS Collaboration. "A portrait of the Higgs boson by the CMS experiment ten years after the discovery". In: *Nature* 607.7917 (July 2022), pp. 60–68. ISSN: 1476-4687. DOI: 10.1038/s41586-022-04892-x. URL: http://dx.doi.org/10.1038/s41586-022-04892-x.

[16] K. G. Begeman. "H I rotation curves of spiral galaxies. I - NGC 3198". In: *Astron. Astrophys.* 223 (1989), pp. 47–60.

[17] SNO Collaboration. "Measurement of the Rate of $\nu_e + d \to p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory". In: *Phys. Rev. Lett.* 87 (7 July 2001), p. 071301. DOI: 10.1103/PhysRevLett.87.071301. URL: https://link.aps.org/doi/10.1103/PhysRevLett.87.071301.

[18] Super-Kamiokande Collaboration. "Evidence for Oscillation of Atmospheric Neutrinos". In: *Phys. Rev. Lett.* 81 (8 Aug. 1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. URL: https://link.aps.org/doi/10.1103/PhysRevLett.81.1562.

[19] Daya Bay Collaboration. "Observation of Electron-Antineutrino Disappearance at Daya Bay". In: *Phys. Rev. Lett.* 108 (17 Apr. 2012), p. 171803. DOI: 10.1103/PhysRevLett.108.171803. URL: https://link.aps.org/doi/10.1103/PhysRevLett.108.171803.

[20] T2K Collaboration. "Evidence of electron neutrino appearance in a muon neutrino beam". In: *Phys. Rev. D* 88 (3 Aug. 2013), p. 032002. DOI: 10.1103/PhysRevD.88.032002. URL: https://link.aps.org/doi/10.1103/PhysRevD.88.032002.

[21] Z. Maki, M. Nakagawa, and S. Sakata. "Remarks on the Unified Model of Elementary Particles". In: *Progress of Theoretical Physics* 28.5 (Nov. 1962), pp. 870–880. ISSN: 0033-068X. DOI: 10.1143/ptp.28.870. URL: http://dx.doi.org/10.1143/PTP.28.870.

[22] R. N. Mohapatra and J. C. Pati. ""Natural" left-right symmetry". In: *Phys. Rev. D* 11 (9 May 1975), pp. 2558–2561. DOI: 10.1103/PhysRevD.11.2558. URL: https://link.aps.org/doi/10.1103/PhysRevD.11.2558.

[23] J. C. Pati and A. Salam. "Lepton number as the fourth "color"". In: *Phys. Rev. D* 10 (1 July 1974), pp. 275–289. DOI: 10.1103/PhysRevD.10.275. URL: https://link.aps.org/doi/10.1103/PhysRevD.10.275.

[24] J. Schechter and J. W. F. Valle. "Neutrino masses in SU(2) $\bigotimes$ U(1) theories". In: *Phys. Rev. D* 22 (9 Nov. 1980), pp. 2227–2235. DOI: 10.1103/PhysRevD.22.2227. URL: https://link.aps.org/doi/10.1103/PhysRevD.22.2227.

[25] V.A. Kuzmin, V.A. Rubakov, and M.E. Shaposhnikov. "On anomalous electroweak baryon-number non-conservation in the early universe". In: *Physics Letters B* 155.1–2 (May 1985), pp. 36–42. ISSN: 0370-2693. DOI: 10.1016/0370-2693(85)91028-7. URL: http://dx.doi.org/10.1016/0370-2693(85)91028-7.

[26] T. Asaka and M. Shaposhnikov. "The νMSM, dark matter and baryon asymmetry of the universe". In: *Physics Letters B* 620.1–2 (July 2005), pp. 17–26. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2005.06.020. URL: http://dx.doi.org/10.1016/j.physletb.2005.06.020.

[27] M. Fukugita and T. Yanagida. "Barygenesis without grand unification". In: *Physics Letters B* 174.1 (June 1986), pp. 45–47. ISSN: 0370-2693. DOI: 10.1016/0370-2693(86)91126-3. URL: http://dx.doi.org/10.1016/0370-2693(86)91126-3.

[28] S. Dodelson and L. M. Widrow. "Sterile neutrinos as dark matter". In: *Phys. Rev. Lett.* 72 (1 Jan. 1994), pp. 17–20. DOI: 10.1103/PhysRevLett.72.17. URL: https://link.aps.org/doi/10.1103/PhysRevLett.72.17.

[29] S. Alekhin et al. "A facility to search for hidden particles at the CERN SPS: the SHiP physics case". In: *Reports on Progress in Physics* 79.12 (Oct. 2016), p. 124201. ISSN: 1361-6633. DOI: 10.1088/0034-4885/79/12/124201. URL: http://dx.doi.org/10.1088/0034-4885/79/12/124201.

[30] J.-L. Tastet, O. Ruchayskiy, and I. Timiryasov. "Reinterpreting the ATLAS bounds on heavy neutral leptons in a realistic neutrino oscillation model". In: *Journal of High Energy Physics* 2021.12 (Dec. 2021). ISSN: 1029-8479. DOI: 10.1007/jhep12(2021)182. URL: http://dx.doi.org/10.1007/JHEP12(2021)182.

[31] I. Esteban et al. "The fate of hints: updated global analysis of three-flavor neutrino oscillations". In: *Journal of High Energy Physics* 2020.9 (Sept. 2020). ISSN: 1029-8479. DOI: 10.1007/jhep09(2020)178. URL: http://dx.doi.org/10.1007/JHEP09(2020)178.

[32] D. A. Trischuk. "Displaced Vertex Search For Heavy Neutral Leptons Using The ATLAS Detector". Presented 19 Jan 2022. University of British Columbia, 2022. URL: https://cds.cern.ch/record/2806047.

[33] M. Gronau, C. N. Leung, and J. L. Rosner. "Extending limits on neutral heavy leptons". In: *Phys. Rev. D* 29 (11 June 1984), pp. 2539–2558. DOI: 10.1103/PhysRevD.29.2539. URL: https://link.aps.org/doi/10.1103/PhysRevD.29.2539.

[34] K. Bondarenko et al. "Phenomenology of GeV-scale heavy neutral leptons". In: *Journal of High Energy Physics* 2018.11 (Nov. 2018). ISSN: 1029-8479. DOI: 10.1007/jhep11(2018)032. URL: http://dx.doi.org/10.1007/JHEP11(2018)032.

[35] L. Evans and P. Bryant. "LHC Machine". In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08001. DOI: 10.1088/1748-0221/3/08/S08001. URL: https://dx.doi.org/10.1088/1748-0221/3/08/S08001.

[36] *The HL-LHC Project Website*. URL: https://hilumilhc.web.cern.ch/content/hl-lhc-project.

[37] O. Aberle et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: 10.23731/CYRM-2020-0010. URL: https://cds.cern.ch/record/2749422.

[38] E. Lopienska. "The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022". In: (2022). General Photo. URL: https://cds.cern.ch/record/2800984.

[39] *Public ATLAS Luminosity Results for Run-2 of the LHC*. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2.

[40] ATLAS Collaboration. "The ATLAS Experiment at the CERN Large Hadron Collider". In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003. URL: https://dx.doi.org/10.1088/1748-0221/3/08/S08003.

[41] J. Pequenao. "Computer generated image of the whole ATLAS detector". 2008. URL: https://cds.cern.ch/record/1095924.

[42] M. Capeans et al. *ATLAS Insertable B-Layer Technical Design Report*. Tech. rep. 2010. URL: https://cds.cern.ch/record/1291633.

[43] ATLAS Collaboration. *ATLAS inner detector: Technical Design Report, 1*. Technical design report. ATLAS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331063.

[44] ATLAS Collaboration. *Alignment of the ATLAS Inner Detector and its Performance in 2012*. Tech. rep. Geneva: CERN, 2014. URL: https://cds.cern.ch/record/1741021.

[45] ATLAS Collaboration. *IBL Efficiency and Single Point Resolution in Collision Events*. Tech. rep. Geneva: CERN, 2016. URL: https://cds.cern.ch/record/2203893.

[46] R. L. Workman et al. "Review of Particle Physics". In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097.

[47] J. Pequenao. "Computer Generated image of the ATLAS calorimeter". 2008. URL: https://cds.cern.ch/record/1095927.

[48] J. Pequenao. "Computer generated image of the ATLAS Muons subsystem". 2008. URL: https://cds.cern.ch/record/1095929.

[49] ATLAS Collaboration. "Operation of the ATLAS trigger system in Run 2". In: *Journal of Instrumentation* 15.10 (Oct. 2020), P10004. DOI: 10.1088/1748-0221/15/10/P10004. URL: https://dx.doi.org/10.1088/1748-0221/15/10/P10004.

[50] J. Pequenao and P. Schaffner. "How ATLAS detects particles: diagram of particle paths in the detector". 2013. URL: https://cds.cern.ch/record/1505342.

[51] ATLAS Collaboration. "Jet reconstruction and performance using particle flow with the ATLAS Detector". In: *Eur. Phys. J. C* 77 (2017), p. 466. DOI: 10.1140/epjc/s10052-017-5031-2. arXiv: 1703.10485 [hep-ex].

[52] *ATLAS Tracking Software Tutorial Website*. URL: https://atlassoftwaredocs.web.cern.ch/trackingTutorial.

[53] R. Frühwirth. "Application of Kalman filtering to track and vertex fitting". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 262.2 (1987), pp. 444–450. ISSN: 0168-9002. DOI: https://doi.org/10.1016/0168-9002(87)90887-4. URL: https://www.sciencedirect.com/science/article/pii/0168900287908874.

[54] W. Waltenberger, R. Frühwirth, and P. Vanlaer. "Adaptive vertex fitting". In: *Journal of Physics G: Nuclear and Particle Physics* 34.12 (Nov. 2007), N343. DOI: 10.1088/0954-3899/34/12/N01. URL: https://dx.doi.org/10.1088/0954-3899/34/12/N01.

[55] ATLAS Collaboration. "Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13\,\text{TeV}$". In: *Eur. Phys. J. C* 81 (2021), p. 578. DOI: 10.1140/epjc/s10052-021-09233-2. arXiv: 2012.00578 [hep-ex].

[56] S. Rettie. "Muon identification and performance in the ATLAS experiment". In: *PoS* DIS2018 (2018), p. 097. DOI: 10.22323/1.316.0097.

[57] ATLAS Collaboration. "Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13\,\text{TeV}$". In: *Eur. Phys. J. C* 79 (2019), p. 639. DOI: 10.1140/epjc/s10052-019-7140-6. arXiv: 1902.04655 [physics.ins-det].

[58] ATLAS Collaboration. *Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung*. Tech. rep. Geneva: CERN, 2012. URL: https://cds.cern.ch/record/1449796.

[59] M. Cacciari, G. P. Salam, and G. Soyez. "The anti-kt jet clustering algorithm". In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. URL: https://dx.doi.org/10.1088/1126-6708/2008/04/063.

[60] ATLAS Collaboration. "Search for Heavy Neutral Leptons in Decays of W Bosons Using a Dilepton Displaced Vertex in $\sqrt{s} = 13$ TeV pp Collisions with the ATLAS Detector". In: *Phys. Rev. Lett.* 131 (6 Aug. 2023), p. 061803. DOI: 10.1103/PhysRevLett.131.061803. URL: https://link.aps.org/doi/10.1103/PhysRevLett.131.061803.

[61] E. Fernández-Martínez et al. "Effective portals to heavy neutral leptons". In: *JHEP* 09 (2023), p. 001. DOI: 10.1007/JHEP09(2023)001. arXiv: 2304.06772 [hep-ph].

[62] S. Agostinelli et al. "Geant4—a simulation toolkit". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: https://doi.org/10.1016/S0168-9002(03)01368-8. URL: https://www.sciencedirect.com/science/article/pii/S0168900203013688.

[63] J. Alwall et al. "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations". In: *JHEP* 07 (2014), p. 079. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph].

[64] C. Degrande et al. "Fully automated precision predictions for heavy neutrino production mechanisms at hadron colliders". In: *Phys. Rev. D* 94 (5 Sept. 2016), p. 053002. DOI: 10.1103/PhysRevD.94.053002. URL: https://link.aps.org/doi/10.1103/PhysRevD.94.053002.

[65] NNPDF Collaboration, Richard D. Ball, et al. "Parton distributions for the LHC run II". In: *JHEP* 04 (2015), p. 040. DOI: 10.1007/JHEP04(2015)040. arXiv: 1410.8849 [hep-ph].

[66] T. Sjöstrand et al. "An introduction to PYTHIA 8.2". In: *Comput. Phys. Commun.* 191 (2015), p. 159. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph].

[67] ATLAS Collaboration. *ATLAS Pythia 8 tunes to* 7 *TeV data.* ATL-PHYS-PUB-2014-021. 2014. URL: https://cds.cern.ch/record/1966419.

[68] NNPDF Collaboration, Richard D. Ball, et al. "Parton distributions with LHC data". In: *Nucl. Phys. B* 867 (2013), p. 244. DOI: 10.1016/j.nuclphysb.2012.10.003. arXiv: 1207.1303 [hep-ph].

[69] S. Frixione, G. Ridolfi, and P. Nason. "A positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction". In: *JHEP* 09 (2007), p. 126. DOI: 10.1088/1126-6708/2007/09/126. arXiv: 0707.3088 [hep-ph].

[70] P. Nason. "A new method for combining NLO QCD with shower Monte Carlo algorithms". In: *JHEP* 11 (2004), p. 040. DOI: 10.1088/1126-6708/2004/11/040. arXiv: hep-ph/0409146.

[71] S. Frixione, P. Nason, and C. Oleari. "Matching NLO QCD computations with parton shower simulations: the POWHEG method". In: *JHEP* 11 (2007), p. 070. DOI: 10.1088/1126-6708/2007/11/070. arXiv: 0709.2092 [hep-ph].

[72] S. Alioli et al. "A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX". In: *JHEP* 06 (2010), p. 043. DOI: 10.1007/JHEP06(2010)043. arXiv: 1002.2581 [hep-ph].

[73] ATLAS Collaboration. *Studies on top-quark Monte Carlo modelling for Top2016.* ATL-PHYS-PUB-2016-020. 2016. URL: https://cds.cern.ch/record/2216168.

[74] D. J. Lange. "The EvtGen particle decay simulation package". In: *Nucl. Instrum. Meth. A* 462 (2001), p. 152. DOI: 10.1016/S0168-9002(01)00089-4.

[75] T. Gleisberg et al. "Event generation with SHERPA 1.1". In: *JHEP* 02 (2009), p. 007. DOI: 10.1088/1126-6708/2009/02/007. arXiv: 0811.4622 [hep-ph].

[76] S. Höche et al. "QCD matrix elements and truncated showers". In: *JHEP* 05 (2009), p. 053. DOI: 10.1088/1126-6708/2009/05/053. arXiv: 0903.1219 [hep-ph].

[77] E. Bothmann et al. "Event generation with Sherpa 2.2". In: *SciPost Phys.* 7.3 (2019), p. 034. DOI: 10.21468/SciPostPhys.7.3.034. arXiv: 1905.09127 [hep-ph].

[78] ATLAS Collaboration. "Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data". In: *JINST* 14 (2019), P12006. DOI: 10.1088/1748-0221/14/12/P12006. arXiv: 1908.00005 [hep-ex].

[79] ATLAS Collaboration. *Neural Network Jet Flavour Tagging with the Upgraded ATLAS Inner Tracker Detector at the High-Luminosity LHC*. Tech. rep. Geneva: CERN, 2022. URL: https://cds.cern.ch/record/2839913.

[80] ATLAS Collaboration. *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*. Tech. rep. Geneva: CERN, 2019. URL: https://cds.cern.ch/record/2669425.

[81] V. Kostyukhin. *VKalVrt - package for vertex reconstruction in ATLAS.* Tech. rep. revised version number 1 submitted on 2003-09-24 11:10:53. Geneva: CERN, 2003. URL: https://cds.cern.ch/record/685551.

[82] ATLAS Collaboration. "Search for long-lived, massive particles in events with displaced vertices and multiple jets in *pp* collisions at $\sqrt{s} = 13\,\text{TeV}$ with the ATLAS detector". In: *JHEP* 06 (2023), p. 200. DOI: 10.1007/JHEP06(2023)200. arXiv: 2301.13866 [hep-ex].

[83] *Monte Carlo Truth Classifier (ATLAS Internal Link)*. URL: https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/MCTruthClassifier.

[84] N. Barlow et al. *Search for long-lived, massive particles decaying into leptons in pp collisions at 8 TeV centre-of-mass energy, with the ATLAS detector at the LHC*. Tech. rep. Geneva: CERN, 2014. URL: https://cds.cern.ch/record/1749971.

[85] ATLAS Collaboration. "Luminosity determination in *pp* collisions at $\sqrt{s} = 13\,\text{TeV}$ using the ATLAS detector at the LHC". In: *Eur. Phys. J. C* 83 (2023), p. 982. DOI: 10.1140/epjc/s10052-023-11747-w. arXiv: 2212.09379 [hep-ex].

[86] ATLAS Collaboration. *Evaluating statistical uncertainties and correlations using the bootstrap method*. Tech. rep. Geneva: CERN, 2021. URL: https://cds.cern.ch/record/2759945.

[87] J. Butterworth et al. "PDF4LHC recommendations for LHC Run II". In: *Journal of Physics G: Nuclear and Particle Physics* 43.2 (Jan. 2016), p. 023001. DOI: 10.1088/0954-3899/43/2/023001. URL: https://dx.doi.org/10.1088/0954-3899/43/2/023001.

[88] K. Cranmer et al. *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*. Tech. rep. New York: New York U., 2012. DOI: 10.17181/CERN-OPEN-2012-016. URL: https://cds.cern.ch/record/1456844.

[89] A. L. Read. "Presentation of search results: the CLs technique". In: *Journal of Physics G: Nuclear and Particle Physics* 28.10 (Sept. 2002), p. 2693. DOI: 10.1088/0954-3899/28/10/313. URL: https://dx.doi.org/10.1088/0954-3899/28/10/313.

[90] G. Cowan et al. "Asymptotic formulae for likelihood-based tests of new physics". In: *Eur. Phys. J. C* 71 (2011), p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [physics.data-an]. Erratum: in: *Eur. Phys. J. C* 73 (2013), p. 2501. DOI: 10.1140/epjc/s10052-013-2501-z.

[91] ATLAS Collaboration. "Measurement of the W-boson mass in pp collisions at $\sqrt{s} = 7\,\text{TeV}$ with the ATLAS detector". In: *Eur. Phys. J. C* 78 (2018), p. 110. DOI: 10.1140/epjc/s10052-017-5475-4. arXiv: 1701.07240 [hep-ex]. Erratum: in: *Eur. Phys. J. C* 78 (2018), p. 898. DOI: 10.1140/epjc/s10052-018-6354-3.

[92] Particle Data Group Collaboration. "Review of Particle Physics". In: *PTEP* 2020.8 (2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.

Appendix A

PERSONAL CONTRIBUTIONS

Contributions to This Analysis

The work presented in this dissertation is a collaborative effort taken up by members of the ATLAS displaced HNL team. As the lead student on this analysis, I have contributed to all the aspects of this search. I have designed, optimized, and validated the secondary vertexing algorithm parameters as described in chapter 3. The algorithm is critical since its the highest point of efficiency loss in the analysis, and ensures that future channel expansions are possible. I have contributed to the background studies and estimation strategies described in detail in chapter 4, including the implementation and optimization of the discriminant variable $\mathcal{S}$ cut, the implementation of the DV lepton isolation (as well as the close-by-effects), and the region definitions. The systematic uncertainties (see section 4.4) considered in this analysis are put in place by me. Specifically, I have contributed independently to the standardized use of LRT muons, including their identification WP calibration and uncertainty evaluation. I designed the standard–LRT muon overlap removal strategy which maximized signal retention efficiency. I contributed to the design of the signal extraction strategy using which I performed the fits to obtain the exclusion contours shown in chapter 5. I am responsible for maintaining all the data formats used in this analysis, starting from DAOD_LLP1, analysis ntuples with systematics, and micro-ntuples which contain information per channel for the six channels considered. Lastly, as the editor of the internal documentation, I ensure that all studies pertaining to the analysis are documented well for streamlined internal reviews and future publications.

Other Contributions to ATLAS

Apart from this analysis, I have had the opportunity to contribute to multiple projects within ATLAS as its member since 2019.

Detector Development: I was stationed at Brookhaven National Lab from August, 2019 to July, 2021 as a key member of the stave assembly team for the ATLAS Inner Tracker upgrade. I assembled the first stave with 28 modules with the latest chip configuration which helped the stave assembly project pass the ATLAS final design review. I was the lead student in the effort to optimize the quality control criteria for the construction of stave cores, coordinating the design and manufacturing between multiple labs in the UK and US. My final contribution to the project was the development of a new metrology system that combined a camera and a laser to yield the first 3D mapping of staves.

Detector Performance: Working on the search for HNLs gave me the opportunity to join a coordinated effort within ATLAS to provide recommendations on how to use displaced muons for LLP analyses. I lead this activity for ATLAS, developing new methods to calibrate the efficiencies of non-standard displaced muons. I work closely with the tracking and the muon combined performance groups for this activity, and my results are used by all analyses which probe signatures using LRT muons.

Physics Analyses: During my time at BNL, I started working on cross-section measurements of the $H \to WW^*$ process in the vector boson fusion production channel. I started with optimizing the hyperparameters of and training the boosted decision trees used in the analysis to discriminate the signal against background. I took up the task to evaluate the impact of theory modeling systematics on the various processes. Eventually, I stepped into the role of the lead student, optimizing the final binning of the results reported in the analysis and liaising with the journal referees to address comments. I wrote the framework to evaluate the correlations between various differential cross-section measurements using a bootstrapping mechanism, which formed the last results to go into the paper making the analysis complete. My involvement in the $H \to WW^*$ group enabled me to contribute to two other analysis in the same final state but different production mechanism and goals. Furthermore, my expertise in the muon combined performance group allowed me to contribute to the $H \to ZZ^* \to 4\ell$ analysis, where I performed studies to understand the impact of the `Loose` identification WP on the analysis. This analysis was the first Run-3 ATLAS paper.

Appendix B

HNL MASS RECONSTRUCTION

This chapter outlines the calculation of the mass of the HNL candidate m_{HNL} which is the main discriminating variable used in the signal region. The calculation is obtained with permission from [60].

The HNL mass (m_{HNL}) can be obtained using energy–momentum conservation in the HNL production ($W \to \mathcal{N}\ell_1$) and decay ($\mathcal{N} \to \ell_2\ell_3\nu$), where ℓ_1 is the prompt lepton and ℓ_2 and ℓ_3 are the charged leptons in the DV. The problem can be summarized with the following equations. Four-momentum conservation in the $\mathcal{N}$ decay gives

$$p^\mu_{\mathcal{N}} = p^\mu_2 + p^\mu_3 + p^\mu_\nu \equiv p^\mu_{23} + p^\mu_\nu. \tag{B.1}$$

Four-momentum conservation in the W decay gives

$$p^\mu_W = p^\mu_1 + p^\mu_{\mathcal{N}} = p^\mu_1 + p^\mu_{23} + p^\mu_\nu. \tag{B.2}$$

The following are defined

$$\begin{aligned}
&p^2_{23} = E^2_{23} - |\vec{p}_{23}|^2 \equiv m^2_{23} \\
&p^{\parallel}_{23} \equiv \vec{p}_{23} \cdot \hat{\boldsymbol{v}} \\
&p^{\perp}_{23} \equiv |\vec{p}_{23} - p^{\parallel}_{23}\hat{\boldsymbol{v}}|
\end{aligned}$$

where m, E, and $|\vec{p}|$ are the mass, energy, and momentum-vector magnitude of the particles indicated by their subscript and $\hat{\boldsymbol{v}}$ is the flight direction of the HNL given by the vector connecting the PV and DV.

The solution to the HNL mass is presented in the coordinate system $k = (\hat{\boldsymbol{x}}', \hat{\boldsymbol{y}}', \hat{\boldsymbol{z}}')$, which is rotated relative to the ATLAS coordinate system, such that the origin of the k-frame is at the PV and the z'-axis points along the flight direction of the HNL. The definition of this coordinate system is

$$\hat{\boldsymbol{z}}' = \hat{\boldsymbol{v}}, \quad \hat{\boldsymbol{x}}' = \frac{\vec{p}_{23} \times \hat{\boldsymbol{z}}'}{|\vec{p}_{23} \times \hat{\boldsymbol{z}}'|}, \quad \hat{\boldsymbol{y}}' = \hat{\boldsymbol{z}}' \times \hat{\boldsymbol{x}}'.$$

The momenta of ℓ_2 and ℓ_3 constrain the components of the neutrino momentum orthogonal to $\vec{p}_{\mathcal{N}}$. This means that energy–momentum conservation in the W and $\mathcal{N}$ decays can be expressed in terms of one unknown variable α, which is the component of neutrino momentum in the $\hat{\boldsymbol{z}}'$ direction. To express eqs. (B.1) and (B.2) in terms of α, the following quantities are defined

$$\vec{p'}_{23} \equiv \vec{q} \tag{B.3}$$

$$\vec{q} = (0,\ |\vec{p}_{23} \times \hat{\boldsymbol{z}}'| \equiv q_\perp,\ \vec{p}_{23} \cdot \hat{\boldsymbol{z}}' \equiv q_z) \tag{B.4}$$

$$\vec{p'}_\nu = (0,\ -q_\perp,\ \alpha) \tag{B.5}$$

$$E'_\nu = \sqrt{q_\perp^2 + \alpha^2}. \tag{B.6}$$

Squaring eq. (B.2) gives

$$p'^2_W = m_W^2 = m_1^2 + m_\nu^2 + m_{23}^2 + 2p'_1 \cdot (p'_{23} + p'_\nu) + 2p'_{23} \cdot p'_\nu \tag{B.7}$$

where

$$p'_1 \cdot (p'_{23} + p'_\nu) = E'_1(E'_{23} + E'_\nu) - p'_{1,z}(q_z + \alpha)$$
$$p'_{23} \cdot p'_\nu = E'_{23}E'_\nu + q_\perp^2 - q_z\alpha.$$

In the energy regime of interest, the charged leptons and neutrino can be treated as massless particles, such that $m_1 = m_\nu = 0$. Rearranging eq. (B.7) to solve for E_ν gives

$$E'_\nu = A + B\alpha \tag{B.8}$$

where

$$A = \frac{(m_W^2 - m_{23}^2)/2 - E_1' E_{23}' + p_{1,z}' q_z - q_\perp^2}{E_1' + E_{23}'}, \quad B = \frac{p_{1,z}' + q_z}{E_1' + E_{23}'}$$

Subtracting eq. (B.8) from eq. (B.6) gives the following quadratic expression in α

$$(B^2 - 1)\alpha^2 + 2AB\alpha + A^2 - q_\perp^2 = 0.$$

The solution for α is therefore

$$\alpha = \frac{-AB \pm \sqrt{(B^2 - 1)q_\perp^2 + A^2}}{(B^2 - 1)}. \tag{B.9}$$

Both solutions for α were studied using simulated HNL events and it was noted that the solution that led to a smaller $|\vec{p}_\mathcal{N}|$ typically led to a value for m_HNL that was closer to the simulated $m_\mathcal{N}$. This solution often corresponded to forward emission of the neutrino with respect to the HNL decay. Therefore, the definition of m_HNL in the analysis uses the solution with the positive radical.

The expression for α in eq. (B.9) depends on m_W. ATLAS has measured the W-boson pole mass to be $M_W = 80.370 \pm 0.019$ GeV [91]. This measurement is combined in Ref. [92] with results from other collider experiments to provide a measurement of the W-boson width, $\Gamma_W = 2.195 \pm 0.083$ GeV. Since the W mass has a width, then if $m_W = M_W$ in eq. (B.9) it is possible that there is no real solution for α. Instead of rejecting these events, m_W is set equal to the median W mass in the kinematically allowed region ($m_{W,\mathrm{med}}$). This ensures that α (and correspondingly m_HNL) always has a real solution.

To define the kinematically allowed region, the minimum W mass that is consistent with the charged-lepton decay products ($m_{W,\mathrm{min}}$) is computed. From eq. (B.7), the mass of the

W boson is given by

$$m_W^2 = m_{23}^2 + 2\left(E_1'E_{23}' + E_\nu'(E_1' + E_{23}') - p_{1,z}'q_z + q_\perp^2 - \alpha(p_{1,z}' + q_z)\right) \tag{B.10}$$

and $m_{W,\min}$ occurs where

$$\frac{d(m_W^2/2)}{d\alpha} = (E_1' + E_{23}')\frac{dE_\nu'}{d\alpha} - (p_{1,z}' + q_z) = 0. \tag{B.11}$$

Using

$$\frac{dE_\nu'}{d\alpha} = \frac{d\sqrt{q_\perp^2 + \alpha^2}}{d\alpha} = \frac{\alpha}{E_\nu'}$$

in eq. (B.11), the chosen value of α that gives the minimum m_W is

$$\alpha = \frac{q_\perp B}{\sqrt{1 - B^2}}. \tag{B.12}$$

Substituting eq. (B.12) into eq. (B.10), the minimum W boson mass is

$$m_{W,\min}^2 = m_{23}^2 + 2\left(E_1'E_{23}' + (E_1' + E_{23}')\sqrt{q_\perp^2 + \frac{q_\perp^2 B^2}{1 - B^2}} - p_{1,z}'q_z + q_\perp^2 - (p_{1,z}' + q_z)\frac{q_\perp B}{\sqrt{1 - B^2}}\right).$$

The cumulative probability for the W boson to have a mass greater than $m_{W,\min}$ is used to find the median of the remaining distribution. The probability density function (f) for m_W^2 satisfies

$$f(m_W^2) \propto \frac{1}{(m_W^2 - M_W^2)^2 + M_W^2\Gamma_W^2}.$$

Therefore, the cumulative distribution function (F) is

$$F(m_W^2) = \frac{1}{\pi}\arctan\left(\frac{m_W^2 - M_W^2}{M_W\Gamma_W}\right) + \frac{1}{2}. \tag{B.13}$$

The midpoint of the allowed kinematic region has a value of

$$F_{\text{med}} = \frac{1 + F(m^2_{W,\text{min}})}{2}$$

Rearranging eq. (B.13) for m^2_W gives

$$m^2_W = M^2_W + \Gamma_W M_W \tan\left(\pi\left[F - \frac{1}{2}\right]\right). \tag{B.14}$$

Substituting $F = F_{\text{med}}$ in eq. (B.14) gives an expression for the median W mass in the kinematically allowed region

$$m^2_{W,\text{med}} = M^2_W + \Gamma_W M_W \tan\left(\pi\left[\frac{1 + F(m^2_{W,\text{min}})}{2} - \frac{1}{2}\right]\right).$$

This value of $m_{W,\text{med}}$ is used in eq. (B.9) to solve for α.

From eq. (B.1) and the definitions in eqs. (B.3) to (B.6), the expression for the HNL mass in terms of α is

$$\begin{aligned} m^2_{\text{HNL}} &= m^2_{23} + 2p'_\nu \cdot p'_{23} \\ &= m^2_{23} + 2E'_{23}\sqrt{q^2_\perp + \alpha^2} + 2q^2_\perp - 2q_z\alpha. \end{aligned} \tag{B.15}$$

Substituting the expression for α in eq. (B.9) into eq. (B.15) gives the solution for the HNL mass.

Appendix C

ASYMPTOTIC APPROXIMATION VERSUS TOYS

The exclusion limits shown in section 5.2 uses an Asymptotic approximation on q_{μ}. Typically, such limits are instead reported using pseudo-experiments, or *toys*, where a large number of fits are performed over a range of injected μ values to then obtain the CL_s value for each μ. For each toy, a pseudo-dataset is created assuming a Poisson distribution in each bin with the mean value equal to the expected number of signal (with a strength equal to the injected μ) plus background events. The signal is then extracted using each such dataset, and the average $\hat{\mu}$ is obtained for that point. This "scan" is used to interpolate and estimate μ_{limit}, along with the corresponding CL_s. For the comparison in this chapter, one hundred pseudo-experiments for the B and S+B hypotheses were performed over a range with one hundred points.

A comparison is presented for two points at the boundary and one in the bulk of the exclusion contour for the 2QDH(IH) model since it presents the most complex fit structure with signal extracted from all regions. The values obtained for this comparison are summarized in table C.1.

m_{HNL}, $c\tau_{\text{HNL}}$=	3 GeV, 100 GeV		10 GeV, 10 GeV		15 GeV, 10 mm	
	Asymptotic	Toys	Asymptotic	Toys	Asymptotic	Toys
Observed 95%	0.16	0.26	0.60	0.62	3.9	4.1
Expected 95%	0.23	0.27	0.53	0.59	3.7	4.1
Expected $\pm 1\sigma$ range	[0.16, 0.33]	[0.25, 0.42]	[0.32, 0.89]	[0.37, 0.88]	[2.2, 6.4]	[2.2, 5.4]
Expected $\pm 2\sigma$ range	[0.11, 0.48]	[0.19, -]	[0.21, 1.46]	[0.29, -]	[1.4, 10.6]	[0.4, 9.1]

Table C.1: Comparisons of the observed and expected limits and the errors on the expected limits obtained using the asymptotic approximation and toys for three representative HNL points in the 2QDH(IH) model. The $+2\sigma$ value is not obtained with toys for some samples and is shown using a dash (-).

The comparison from these points shows that the central values are compatible within errors for the two methods of upper limit extraction. Furthermore, the range of the errors are similar between the two methods, which depicts that the asymptotic approximation is adept at estimating not just the central values accurately but also the error on them.