

Titel der Arbeit:

Extending the limits in the hunt for long-lived heavy neutral leptons with the ATLAS experiment at the Large Hadron Collider at CERN

Dissertation

zur Erlangung des akademischen Grades

doctor rerum naturalium

(Dr. rer. nat.)

im Fach Physik: Spezialisierung:

Experimentalphysik

eingereicht an der

Mathematisch-Naturwissenschaftlichen Fakultät der Humboldt-Universität zu Berlin

von

M.Sc. Christian Appelt

Präsidentin der Humboldt-Universität zu Berlin

Prof. Dr. Julia von Blumenthal

Dekanin der Mathematisch-Naturwissenschaftlichen Fakultät

Prof. Dr. Caren Tischendorf

Gutachter:

1. Prof. Dr. Heiko Lacker
2. Prof. Dr. Steven Worm
3. Prof. Dr. Ulrich Husemann

Tag der mündlichen Prüfung: 02.05.2024




Extending the limits in the hunt for long-lived heavy neutral leptons

with the ATLAS experiment at the Large Hadron Collider at CERN

by

Christian Appelt



Directed by:

Prof. Dr. Heiko Lacker

ABSTRACT

A search for a long-lived, heavy neutral lepton ($\mathcal{N}$) in 139 fb^{-1} of $\sqrt{s} = 13 \text{ TeV}$ pp collision data collected between 2015 and 2018 by the ATLAS detector at the Large Hadron Collider is presented. The $\mathcal{N}$ is assumed to be produced via $W \rightarrow \mathcal{N}\mu$ or $W \rightarrow \mathcal{N}e$ and decays into two charged leptons and a neutrino, forming a unique displaced vertex signature. A novel technique is used to reconstruct the $\mathcal{N}$ mass that discriminates between signal and background. No signal is observed, and limits are set on the squared mixing parameters of the $\mathcal{N}$ with the left-handed neutrino states for the $\mathcal{N}$ mass range $3 \text{ GeV} < m_{\mathcal{N}} < 15 \text{ GeV}$. For the first time, limits are given for both single-flavour and multi-flavour mixing scenarios motivated by neutrino-flavour oscillation results for both the normal and inverted neutrino-mass hierarchies. In addition to the concluded search for $\mathcal{N}$, ongoing efforts are outlined that include 66 fb^{-1} of $\sqrt{s} = 13.6 \text{ TeV}$ pp collision data collected between 2022 and 2023 by the ATLAS detector and semileptonic $\mathcal{N}$ decays into a charged lepton and a charged pion.

ZUSAMMENFASSUNG

Eine Suche nach langlebigen, schweren neutralen Leptonen ($\mathcal{N}$) in 139 fb^{-1} mit $\sqrt{s} = 13 \text{ TeV}$ pp Kollisionsdaten der Jahre 2015-2018 mit Hilfe des ATLAS-Detektors am Large Hadron Collider ist präsentiert. Die $\mathcal{N}$ werden in $W \rightarrow \mathcal{N}\mu$ oder $W \rightarrow \mathcal{N}e$ produziert und zerfallen in zwei geladene Leptonen und ein Neutrino. Die Zerfallsprodukte des $\mathcal{N}$ bilden eine einzigartige versetzte Vertex-Signatur. Die $\mathcal{N}$ Masse wird mit Hilfe einer neuartigen Technik rekonstruiert und genutzt, um zwischen Signal- und Untergrundereignissen zu differenzieren. Es wurden keine Signalereignisse beobachtet, Sensitivitätsgrenzen werden auf die quadrierten Mischungsparameter zwischen rechtshändigen und linkshändigen Neutrinos für den $\mathcal{N}$ Massenbereich von $3 \text{ GeV} < m_{\mathcal{N}} < 15 \text{ GeV}$ gesetzt. Zum ersten Mal werden diese Sensitivitätsgrenzen für Kopplungen für einen Neutrino-Flavour, aber auch für Kopplungen für mehrere Neutrino-Flavour-Mischungsverhältnisse angegeben, die mit den gemessenen Neutrino-Oszillationsdaten für die invertierte und normale Massenhierarchie übereinstimmen. Zusätzlich zu der bereits abgeschlossenen Suche nach $\mathcal{N}$ werden laufende Aktivitäten skizziert, die 66 fb^{-1} und $\sqrt{s} = 13.6 \text{ TeV}$ pp Kollisionsdaten der Jahre 2022-2023 des ATLAS-Detektors und semi-leptonische $\mathcal{N}$ Zerfälle in ein geladenes Lepton und ein geladenes Pion einbeziehen.

CONTENTS

1	INTRODUCTION	1
2	PREFACE	3
I	MOTIVATION AND EXPERIMENT	9
3	AN EXTENSION TO THE STANDARD MODEL	11
3.1	Standard model of particle physics: a brief summary	11
3.1.1	Constituents and forces	11
3.1.2	Input parameter	11
3.1.3	Symmetries	13
3.1.4	Mass	13
3.1.5	Successes and limitations	14
3.2	Heavy neutral lepton models	16
3.2.1	Overview	16
3.2.2	Type I Seesaw mechanism	19
3.2.3	A Quasi-Dirac pair of heavy neutral leptons	21
4	THE EXPERIMENT	27
4.1	Large Hadron Collider	27
4.2	ATLAS detector	30
4.2.1	Inner detector	31
4.2.2	Calorimeters	33
4.2.3	Muon spectrometer	35
4.2.4	Trigger and data aquisition	35
4.3	Physics object reconstruction	36
4.3.1	Tracks	36
4.3.2	Primary vertices, pileup, and luminosity	41
4.3.3	Muons	42
4.3.4	Electrons	44
4.4	Special long-lived particle reconstruction	49
4.4.1	Large-radius tracking	49
4.4.2	Displaced vertex reconstruction	50
4.5	Simulation	53

II	DATA ANALYSIS	59
5	SEARCH FOR LONG-LIVED HEAVY NEUTRAL LEPTONS	61
5.1	Sensitivity	61
5.2	Signature	65
5.3	Signal models	67
5.4	Samples	71
5.4.1	Signal samples	71
5.4.2	Data samples	73
5.4.3	What about displaced taus?	73
5.4.4	Signal kinematics	74
5.4.5	Additional weights on signal samples	77
5.5	Analysis framework	81
5.5.1	Custom secondary vertex reconstruction	82
5.6	Event selection	84
5.6.1	Object definition	84
5.6.2	Preselection	86
5.6.3	Signal region selection	88
5.6.4	Signal region selection efficiency and cutflow	93
5.6.5	Validation region and control region selection	97
5.7	Background estimation	98
5.7.1	Metastable particle decays	100
5.7.2	Random track crossing background	102
5.7.3	Background validation	105
5.8	Discriminating variable	107
5.9	Systematic uncertainties	111
5.9.1	Signal systematic uncertainties	111
5.9.2	Background systematic uncertainties	117
5.10	Statistical interpretation	122
5.10.1	CLs method	122
5.10.2	Channel combinations and maximum likelihood fit	126
5.10.3	Exclusion limits	132
6	ONGOING SEARCH FOR LONG-LIVED HEAVY NEUTRAL LEPTONS	141
6.1	Improvements in large radius tracking	142
6.2	A new analysis channel	142
6.3	Data-driven background estimate	146
6.4	Muon spectrometer vertexing	146
6.5	Data and MC simulation samples	147

III	FINAL WORDS	151
7	SUMMARY	153
8	OUTLOOK	155
A	APPENDIX	161
A.1	Likelihood and maximum likelihood method	161
A.1.1	Probability density functions	161
A.1.2	Likelihood	162
A.1.3	Maximum likelihood method	162
A.2	HNL mass calculation	163
A.3	Truth reweighting validations	167
A.4	Signal selection efficiencies	170
A.5	Background uncertainties in final binning	172
A.6	Technical details of the DHNLLimits framework	177
A.7	Impact of systematic uncertainties on the expected 95% confidence limit	178
	BIBLIOGRAPHY	181

LIST OF FIGURES

3.1	Constituents of the standard model of particle physics.	12
3.2	Summary of the fermion masses split by generation.	14
3.3	Summary of SM cross-section measurements performed by ATLAS.	15
3.4	Diagrams contributing to neutrino masses in the three Seesaw scenarios at tree level.	19
3.5	Schematic illustration of the particle content of the neutrino minimal SM (ν MSM).	21
3.6	The range of the total coupling $U^2 = U_{\text{tot}} ^2$ consistent with both the Seesaw mechanism and leptogenesis as a function of the average HNL mass $M = M_N$ for the case of two HNLs.	25
4.1	Overview of the LHC and its underground location in Switzerland and France.	28
4.2	Overview of the LHC and the CERN accelerator complex.	29
4.3	Runs and technical shutdowns of the LHC and High Luminosity (HL)-LHC.	29
4.4	Schematic of the ATLAS detector.	30
4.5	Schematic illustration of the inner detector.	31
4.6	The ATLAS inner tracking detector layout.	32
4.7	Schematic illustration of the different calorimeter systems in ATLAS.	33
4.8	Schematic illustration of the different parts of the muon system in ATLAS.	34
4.9	Overview of the trigger and data acquisition system of the ATLAS detector.	36
4.10	Schematic overview of particle candidates and where they are reconstructed in the ATLAS detector.	37
4.11	Schematic illustration defining the track parameters used in ATLAS.	37
4.12	Illustration of single-particle clusters on a pixel sensor and a merged pixel cluster.	38
4.13	A sample event from a $t\bar{t}$ MC simulation showing the two possible TRT hit associations.	39
4.14	Illustrative event display without dimensions of multiple interactions along the beamline (pile-up).	41
4.15	Luminosity weighted distribution of the mean number of interactions per bunch crossing μ for the 2015-2018 pp collision data and cumulative luminosity as a function of time.	42
4.16	Muon reconstruction and identification efficiencies.	45
4.17	Schematic illustration of the path of an electron through the detector.	46

4.18	The transformed LH-based identification discriminant d'_L for reconstructed electron candidates.	46
4.19	Identification efficiencies of electrons.	47
4.20	Inclusive track reconstruction efficiency for displaced charged particles produced by the decay of long-lived signal particles.	50
4.21	Displaced vertex reconstruction scheme.	51
4.22	Reconstruction efficiency and the difference between the true (MC event generator level) and reconstructed decay position for different long-lived particle decays.	52
4.23	Sketch of a hadron-hadron collision as simulated by a Monte Carlo event generator.	54
4.24	Example of MMHT2014 NNLO parton distribution functions (PDFs).	55
5.1	Constraints on $ U_e ^2$ as a function of HNL mass $m_{\mathcal{N}}$ (electron coupling only, Majorana HNLs).	62
5.2	Constraints on $ U_\mu ^2$ as a function of HNL mass $m_{\mathcal{N}}$ (muon coupling only, Majorana HNLs).	63
5.3	Constraints on $ U_\tau ^2$ as a function of HNL mass $m_{\mathcal{N}}$ (tau coupling only, Majorana HNLs).	63
5.4	Feynman diagrams of the leptonic signature considered in the search.	64
5.5	Schematic illustration of the signature in the ATLAS detector we define as <i>signal</i>	66
5.6	Feynman diagram for the HNL decay via an off-shell Z^* boson for the LNC case.	67
5.7	HNL coupling ratios and allowed regions according to recent neutrino oscillation data.	68
5.8	Schematic illustration of the possible channel combinations for the eight models considered in this work.	69
5.9	HNL coupling constant as a function of HNL mass for three different mean decay lengths and the 1SFH(e) model. Branching ratios (BR) for the four models in the Majorana-limit as a function of HNL mass for all considered HNL decay channels.	72
5.10	Truth W boson kinematics.	74
5.11	The distributions of the true transverse momentum (p_T) for the three charged leptons involved in the Heavy Neutral Lepton (HNL) decay.	75
5.12	Truth DV mass and tri-lepton mass ($m_{\ell\ell\ell}$) distributions.	76
5.13	Truth distribution for the radius of the displaced vertex.	76
5.14	Truth HNL kinematics.	78
5.15	Calculated event weights for $w_{\text{CC+NC}}$ and w_{CC}	78
5.16	Truth reweighting validation p_T distributions for the four leptons between MADGRAPH5_AMC@NLO and PYTHIA.	80
5.17	Schematic illustration of the analysis framework and reconstruction chain.	81
5.18	Impact parameter $ d_0 $ distributions of the displaced vertex tracks for two signal MC samples.	82
5.19	The total displaced vertex reconstruction efficiency as a function of displaced vertex radius.	83

5.20	Two-dimensional histograms of the calibrated displaced electron transverse momentum $p_T^{\text{calibrated-electron}}$ as a function of the transverse momentum of the corresponding track $p_T^{\text{track-only}}$	86
5.21	Signal MC distributions for the separation between two displaced muons.	89
5.22	Two-dimensional maps of the observed vertex density in regions vetoed by the material map.	91
5.23	The invariant mass between the prompt lepton and same-flavor lepton with opposite-charge in the displaced vertex for same-sign (SS) displaced vertices in data.	92
5.24	The event selection efficiency for each mass-lifetime point for the analysis channels $\mu-\mu\mu$ and $\mu-\mu e$	93
5.25	Feynman diagram for the HNL production from a Z boson without a prompt lepton that features a two-lepton opposite-sign displaced vertex.	98
5.26	Schematic illustration of possible background sources in this work.	99
5.27	A selection of the SM particle spectrum is shown as a function of mass and proper lifetime.	100
5.28	Two-dimensional histograms of the displaced vertex mass as a function of the displaced vertex radius of opposite-sign displaced vertices in the inverted prompt lepton validation region.	101
5.29	The difference between the number of opposite-sign and same-sign displaced vertices for the validation region in the displaced vertex mass versus displaced radius plane.	103
5.30	The displaced vertex mass distributions in the inverted prompt lepton validation region for opposite-sign and same-sign displaced vertices after applying the heavy flavour decay veto.	104
5.31	Schematic illustration of the objects we shuffle from different events to re-pair and create new artificial events that form the data-driven estimate for the random track crossing background.	105
5.32	HNL mass distributions for the data-driven estimates for the random track crossing background, the observed number of same-sign displaced vertex events, and the observed number of opposite-sign displaced vertex events.	106
5.33	HNL mass distributions of the original same-sign events that serve as a source for the prompt leptons for shuffling and the HNL mass distributions from the shuffled events.	108
5.34	Three-lepton mass distributions of the original same-sign events that serve as a source for the prompt leptons for shuffling and the three lepton mass distributions from the shuffled events.	109
5.35	The reconstructed HNL mass distributions in the signal and control regions for the shuffled background estimate, three different HNL signal MC samples, and the observed opposite-sign data for all six analysis channels.	110
5.36	Reduction of selection efficiency at the displaced vertex type cut after removing a random collection of tracks following the central performance groups recommendations.	113
5.37	The tracking uncertainty in % for the different analysis channels.	114
5.38	Comparison of $K_S^0 \rightarrow \pi^+\pi^-$ event yields in the validation region with MC simulations produced with PYTHIA 8.186 in bins of p_T and r_{DV}	115

5.39	The track d_0 for displaced leptons (muons and electrons) with and without the lepton quality selection applied.	118
5.40	Schematic illustration of the shuffling procedures applied to derive the two systematic uncertainties that arise from the data-driven background estimation.	120
5.41	The HNL mass distributions in the signal and control region for the data-driven background estimate, the observed same sign displaced vertices together with the HNL mass distribution used to estimate the displaced-vertex systematic, and the distribution to estimate the prompt-lepton systematic for all six analysis channels.	121
5.42	Pre-fit and post-fit HNL mass histograms of $m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm HNL signal samples in the two quasi-Dirac normal hierarchy (2QDH (NH)) HNL model in the Majorana limit (including LNC and LNV decays) and the data-driven background estimate for signal and control regions (SR, CR) and all contributing analysis channels.	125
5.43	The ranking plot generated by TREXFITTER that lists the 15 most impactful nuisance parameters Θ ordered by their post-fit impact on the signal strength μ	128
5.44	Single-bin HNL mass histogram of the data-driven shuffled background estimate in the final signal region and control region binning for the $\mu - \mu\mu$ analysis channel.	129
5.45	Listing of the pre-fit histograms for the eight HNL models in Dirac and Majorana-limit. . .	130
5.46	Listing of the post-fit histograms for the eight HNL models in Dirac and Majorana-limit. .	131
5.47	Different confidence levels (CL_b : background only, CL_{s+b} : signal plus background, and $CL_s = CL_{s+b}/CL_b$) as a function of signal strength $\mu = \sigma/\sigma_{theory}$ for the 2QDH (NH) model in Majorana-limit.	134
5.48	Coupling constant versus mass grid for the 2QDH (NH) model in the Majorana-limit. . .	135
5.49	The cubic spline interpolation of the associated signal strength values to each tested point in the HNL coupling constant versus mass plane for the 2QDH (NH) model in the Majorana-limit.	136
5.50	Combination of cubic spline interpolations to obtain the expected and observed exclusion contours in the HNL coupling constant versus mass plane.	136
5.51	Listing of the observed and expected 95% CL contours for the eight different HNL models in the coupling versus mass plane.	137
5.52	The observed and expected 95% CL contours for the ISFH(μ) model in the Dirac-limit compared to the DELPHI limit on the total coupling $ U ^2$ and the limit previously set by ATLAS in the first displaced HNL search.	138
5.53	Summary of the observed 95% CL contours for the eight different HNL models.	139
6.1	Title slide of the analysis kickoff meeting in March 2022 that marks the start of the ongoing search for long-lived heavy neutral leptons.	141

6.2	The processing time taken per event to reconstruct the same data events and event size of the LRT reconstruction output versus average pileup, comparing the legacy LRT (R21) and updated LRT (R24) reconstruction software.	143
6.3	Schematic illustration of the analysis framework and reconstruction chain for the R24 analysis.	144
6.4	Feynman diagrams of the two-body semileptonic signature currently tested in the ongoing search.	145
6.5	Branching ratios (BR) for the 2QDH(NH) model as a function of HNL mass for the leptonic and two-body semileptonic decay channels.	145
6.6	Schematic illustration of the idea to reconstruct displaced vertices using only information from the muon spectrometer.	147
6.7	Luminosity-weighted distributions for the LHC Run 1-3 of the mean number of interactions per bunch crossing. Cumulative luminosity versus time delivered to (green) and recorded by ATLAS (yellow) during stable beams for pp collisions at 13.6 TeV centre-of-mass energy in LHC Run 3 (2022-2023).	148
8.1	A schematic map showing a possible location for the Future Circular Collider at CERN. .	156
8.2	Constraints from past (shaded grey areas) and proposed future experiments (coloured lines) on the HNL coupling to muons only $ U_\mu ^2$ as a function of HNL mass $m_{\mathcal{N}}$	157
A.1	Truth reweighting validation p_T distributions for the four leptons ($l_\alpha, l_\beta, l_\gamma, \nu_\gamma$) between MADGRAPH5_AMC@NLO and PYTHIA for a $\mu-\mu\mu$ ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) sample (LNV).	167
A.2	Truth reweighting validation p_T distributions for the four leptons ($l_\alpha, l_\beta, l_\gamma, \nu_\gamma$) between MADGRAPH5_AMC@NLO and PYTHIA for a $\mu-\mu e$ ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) sample (LNV).	168
A.3	Truth reweighting validation p_T distributions for the four leptons ($l_\alpha, l_\beta, l_\gamma, \nu_\gamma$) between MADGRAPH5_AMC@NLO and PYTHIA for a $\mu-\mu e$ ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) sample (LNC).	169
A.4	The event selection efficiency for each mass-lifetime point for the analysis channels $e-ee$, $e-e\mu$, $e-\mu\mu$, and $\mu-ee$	171
A.5	Single-bin HNL mass histogram of the data-driven shuffled background estimate and its displaced vertex uncertainty in the final signal region binning.	173
A.6	Single-bin HNL mass histogram of the data-driven shuffled background estimate and its displaced vertex uncertainty in the final control region binning.	174
A.7	Single-bin HNL mass histogram of the data-driven shuffled background estimate and its prompt lepton uncertainty in the final control region binning.	175

A.8	Single-bin HNL mass histogram of the data-driven shuffled background estimate its prompt lepton uncertainty in the final signal region binning.	176
A.9	Schematic illustration of the DHNLLimits framework.	178
A.10	The expected 95% CL contours for the eight different models in Dirac-limit and Majorana-limit including and excluding the systematic uncertainties.	179

LIST OF TABLES

4.1	Listing of the main differences between standard and large-radius tracking.	49
5.1	Coupling ratios x_α for the different benchmarks.	71
5.2	Prefactors that need to be added to the final cross-section.	71
5.3	Listing of the DRAW HNL filter criteria.	87
5.4	Description of the selection criteria used in the signal region.	89
5.5	Cutflow for six simulated signal channels showing the weighted number of expected events based on the single-flavour mixing model in the Majorana limit.	95
5.6	Cutflow for the six channels in data showing the number of events passing each successive signal selection for Majorana HNLs.	96
5.7	Description of the selection criteria used in the inverted prompt lepton validation region.	97
5.8	KS test probabilities to test the randomness of prompt lepton and same-sign displaced vertex pairing.	107
5.9	Summary of the systematic uncertainties impacting reconstruction efficiencies.	112
5.10	Summary of the relative systematic uncertainties in the data-driven background estimate.	112
5.11	Definitions of pass-fail combinations with and without large-radius tracking (LRT) enabled.	116
5.12	Post-fit values of the normalisation factors μ_b and the signal strength μ used in the fit for the eight different HNL signal models in Dirac-limit (DL) and Majorana-limit (ML) for ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm).	124

1 INTRODUCTION

This thesis presents a part of the search campaign in the ATLAS experiment at the Large Hadron Collider at CERN for hypothetical particles, so-called heavy neutral leptons, that could give rise to several phenomena not described by the standard model of elementary particle physics.

The standard model of elementary particle physics is the most successful collection of theories we have to date and to our current understanding to describe, predict, and explain the fundamental building blocks of nature and their interactions. Although it is far from perfect and far from being the one ultimate theory that explains everything, the standard model predicts processes observed at the Large Hadron Collider and beyond up to a remarkable level of accuracy.

However, a few observations point to physics beyond the standard model. Studies in astrophysics and cosmology revealed that additional matter of unknown type, so-called dark matter, is needed to describe gravitational effects in the rotation of galaxies, effects seen in gravitational lensing, the large-scale structure of the observable universe, and cosmic microwave background anisotropies. Cosmologists say that approximately 25% of our universe is composed of dark matter [1, 2]. A second fact that the standard model fails to explain is the abundance of matter over antimatter that remained after the Big Bang, also referred to as baryon asymmetry. Finally, the observation of neutrino oscillations [3–5] led to the conclusion that neutrinos have a mass. The observation of neutrino oscillations is the first confirmed evidence for the existence of physics beyond the standard model that can be measured in a laboratory on Earth [6].

All those observations call for an extension of the standard model that explains the massiveness of neutrinos and the abundance of matter in the universe and provides a dark matter candidate. A possible explanation could be the existence of new, massive, right-handed Majorana neutrinos, also called heavy neutral leptons, that complement the left-handed standard model neutrinos [7, 8]. The extended standard model, including heavy neutral leptons, would remain consistent and renormalisable up to the Planck scale [9].

In the ATLAS experiment, we are not sensitive to heavy neutral leptons that could be dark matter candidates because they are too light [7, 8]. Still, we can search for those heavy neutral leptons that could give rise to neutrino masses and the baryon asymmetry in the universe.

The first search presented in this work uses proton-proton collision data recorded between 2015 and 2018 with a centre-of-mass energy of 13 TeV and an integrated luminosity of 139 fb^{-1} . The heavy neutral leptons considered in the first part of this work are assumed to be produced in leptonic W boson decays and decay themselves into two charged leptons and a neutrino. The small coupling values of the heavy neutral leptons

to the standard model neutrinos that are probed in this work make them long-lived, which requires dedicated techniques to reconstruct tracks far away from the beamline and secondary vertices. In the ATLAS detector, we can select events containing a reconstructed charged lepton from the W boson decay. In addition, events are required to contain a significantly displaced vertex with respect to the beamline that consists of precisely two oppositely charged leptons from the heavy neutral lepton decay. Novel techniques are used to determine the data-driven background estimate and to reconstruct the heavy neutral lepton mass that discriminates between signal and background.

No heavy neutral lepton signal is observed. Therefore, exclusion limits are set on the coupling between left-handed and right-handed neutrinos for a heavy neutral lepton mass range of 3-15 GeV. For the first time in the history of collider searches, limits are given for both single-flavour and multi-flavour mixing scenarios motivated by neutrino-flavour oscillation results for both the normal and inverted neutrino-mass hierarchies. The results have been published in 2023 [10].

The second search for long-lived heavy neutral leptons that is ongoing at the time of writing this thesis is outlined, and new ideas are presented to improve the sensitivity. A new analysis channel that includes the two-body semileptonic heavy neutral lepton decay into a charged lepton and a charged pion, an improved background estimation strategy, and the inclusion of proton-proton collision data recorded between 2022 and 2023 with a centre-of-mass energy of 13.6 TeV and an integrated luminosity of 66 fb^{-1} are discussed. In addition, new vertex reconstruction techniques involving muon spectrometer tracks are explored to extend the sensitivity to heavy neutral leptons that decay outside the inner tracker.

This thesis is structured as follows: After a brief preface describing the ATLAS collaboration and my contributions in chapter 2, the standard model of particle physics and its proposed extension by heavy neutral leptons is explained in chapter 3. The following outline of the experiment in chapter 4 summarises the Large Hadron Collider, the ATLAS detector, the physics objects reconstruction such as tracks, muons, electrons, large-radius tracks and secondary vertices, and the principles of hadron collision simulation at the Large Hadron Collider. The main part of this work, the published search for long-lived heavy neutral leptons in the 2015-2018 dataset, is discussed in chapter 5. Chapter 6 sketches the ongoing efforts to search for long-lived heavy neutral leptons, including the 2022-2023 dataset. Chapters 7 and 8 conclude this work and provide a general outlook. The appendix in chapter A lists additional material and figures that will be referred to in the main text.

2 PREFACE

"ATLAS is one of the largest collaborative efforts ever attempted in science" [11]. The ATLAS collaboration has about 6000 members and 3000 scientific authors from 182 institutions in 42 countries. To become a scientific author, each member of the ATLAS collaboration must dedicate half of their time for a year to a service task that benefits the collaboration. My service task was the development of a web application related to the production database of the ATLAS inner tracker (ITk) upgrade that will be installed after the Large Hadron Collider (LHC) Run3 is finished in 2025 (documented in an internal note, ref. [12]). The ATLAS ITk upgrade splits into the pixel [13] and silicon strip [14] detector. My service task was related to the components of the strip detector. The production and acceptance testing of the ITk components is spread over many institutes worldwide, which requires dedicated bookkeeping in a database to track and grade components in every production step. The web application I developed summarises all acceptance test results of a given component to make them easily digestible. I continue to be involved in ITk activities and related to the further development of the web application as the local supervisor for Sadia Marium (PhD student, DESY) and the technical supervisor for Scott Williams (PhD student, Melbourne) for their service tasks.

The LHC Run2 search for long-lived heavy neutral leptons is the main part of this thesis. The search uses proton-proton collision data recorded between 2015 and 2018 with a centre of mass energy of 13 TeV and an integrated luminosity of 139 fb^{-1} and was performed in close collaboration with the team members Dave Casper (University of California (UCI)), Matthias Danninger (Simon Fraser, analysis contact), Sean Galantzan (Tel Aviv), Sergio Grancagnolo (Humboldt University (HU), analysis contact), Heiko Lacker (HU), Alison Lister (University of British Columbia (UBC)), Margaret Lutz (at the time Tel Aviv, analysis contact), Robin Newhouse (UBC), Abi Soffer (Tel Aviv), and Dominique Trischuk (at the time UBC). Oleg Ruchayskiy (Niels Bohr Institute (NBI)) and Jean-Loup Tastet (NBI) contributed to the analysis as short-term associates (external collaborators with temporary authorship of an ATLAS publication). They gave advice related to the heavy neutral lepton theory and model interpretation.

Margaret Lutz and Abi Soffer are the contact editors for the work published in PRL [10]. I created a significant amount of figures for the PRL paper, as well as the final exclusion limit plot in the results part. I also contributed to the writing process. Robin Newhouse contributed significantly to the analysis framework, especially the Ntuple-making code (DHNLAlgorithm), conducted studies related to the DRAW filter, estimated the systematic uncertainties, produced the Ntuples on the Vancouver cluster, and worked with me on the HEP data entry for the analysis. Dominique Trischuk contributed significantly to the second part

of the analysis framework (DHNLLNtupleAnalysis), which performs the displaced vertex selections and produces the micro-Ntuples used as input for the statistical data analysis. Furthermore, Dominique Trischuk estimated and validated the data-driven random track crossing background estimate and the corresponding systematic uncertainties. Dominique Trischuk and Robin Newhouse provided the data as micro-Ntuples I used for the statistical analysis. The framework I wrote is the third part of the analysis algorithm chain: the DHNLLimits framework. The framework involves the final signal and control region selection, the analysis channel combination, the maximum likelihood fit and the exclusion limit determination for each signal sample, and the final interpolation of the exclusion contours in the HNL coupling versus mass plane. I studied the impact of the different systematic and statistical uncertainties on the maximum likelihood fit and the final expected exclusion limits. The DHNLLimits framework creates all summary/results figures and the data tables/figures required for the HEP data entry of the analysis [15]. The low number of expected and observed events in the final signal region selection requires the use of pseudo experiments to create the test-statistic distributions that are used to determine the exclusion limits. Those pseudo experiments are computationally very expensive, which requires the parallelisation of the DHNLLimits framework and the use of several thousand working nodes on a batch system. The scripts that perform this parallelised computation are part of the DHNLLimits framework. Margaret Lutz, Robin Newhouse, Dominique Trischuk and I wrote the significant parts of the ATLAS-internal documentation (INT note, ref. [16]). Dominique Trischuk's PhD thesis [17], Robin Newhouse's PhD thesis [18], and my work presented here have similar content since most of the figures and studies were performed collaboratively, as described above.

After completing the LHC Run2 search for heavy neutral leptons, I defined the analysis strategy at the kick-off meeting for the new search with new ideas and additional analysis channels. The new analysis is conducted in the new ATLAS software release (R24) that involves optimised reconstruction algorithms. Furthermore, we plan to include a part of the LHC Run3 proton-proton collision data recorded between 2022 and 2023 with a centre of mass energy of 13.6 TeV and an integrated luminosity of 66 fb^{-1} and the LHC Run2 dataset. The ongoing analysis is a collaborative effort of Sagar Addepalli (Brandeis), Marzieh Bahmani (HU), Jackson Burzynski (Simon Fraser), Carlo Dallapiccola (Massachusetts), Matthias Danninger (Simon Fraser), Guglielmo Frattari (Brandeis), Maximilian Golbirsch-Kolb (CERN), Alessandro Guida (HU), Audrey Kvam (Massachusetts), Heiko Lacker (HU), Margaret Lutz (CERN), Abi Soffer (Tel Aviv), Dominique Trischuk (Brandeis), and Makayla Vessella (Massachusetts). Audrey Kvam and I are the analysis contacts for the ongoing search for long-lived heavy neutral leptons and manage the analysis team. To get the analysis started, I moved the whole analysis framework (DHNLLAlgorithm, DHNLLNtupleanalysis and DHNLLimits) to the new software release that brought significant changes to the data structure. As one of the first analysis groups in ATLAS working on the new software release, this was a pioneering effort. The LHC Run2 (+Run3) analysis is ongoing at the time of writing this thesis. No results are included, but the improvements compared to the LHC Run2 search and new ideas are outlined in this thesis.

So far, at HU, I have co-supervised six Bachelor's students (Daniel Stefaniak, Ida Wösthreinrich, Leon Buchholz, Larissa Prochnow, Tillmann Cebullat, and Theo Oldenburg) and two Master's students (Julius Ehram and The Anh Tran). For all of them, I suggested thesis topics related to the heavy neutral lepton analysis to Heiko Lacker, and I was responsible for the technical and daily supervision of their work.

ACKNOWLEDGEMENTS

First, I want to thank Heiko Lacker for his helpful advice and guidance during my work as a PhD candidate in his group and for his insightful comments on the draft version of this work.

Secondly, I would like to express my gratitude to my colleagues, who made this work possible. I thank the analysis contacts Matthias Danninger, Sergio Grancagnolo, and Margaret Lutz for their management, organisation, and guidance. I want to thank Dominique Trischuk and Robin Newhouse, who also graduated from this search and made this work a great team effort; it was lovely to seamlessly work with you guys across time zones to meet all the deadlines on time. I want to thank Oleg Ruchayskiy and Jean-Loup Tastet for their priceless advice on HNL theory matters and for the numerous insightful discussions during the analysis meetings that enabled us to include the real physics HNL model interpretations into our work as the first collider search. I want to thank Alison Lister for the guidance and discussions on the statistical interpretation of this work. Finally, I thank Abi Soffer for additional guidance and the little bit of extra motivation to write up fast by giving me a perspective in physics as a PostDoc at CERN.

Thirdly, I thank my colleagues and former colleagues here at Humboldt University: Ferdinand Schenk, Michele Michetti, Christian Scharf, Marzieh Bahmani, Alessandro Guida, Andrew Conaboy, Anupama Reghunath, Yulia Jagodzinski, Cigdem Issever, Thomas Lohse, Steven Worm, and many more for all the delightful social gatherings, the lunches, and sometimes listening to my complaints.

Fourth, thank you, Bruce Gallop, for guiding me through my qualification task to become an ATLAS author.

Finally, I'd like to thank Audrey Kvam for being such a fabulous co-analysis contact in the ongoing search for heavy neutral leptons.

PART I

MOTIVATION AND EXPERIMENT

3 AN EXTENSION TO THE STANDARD MODEL

3.1 STANDARD MODEL OF PARTICLE PHYSICS: A BRIEF SUMMARY

3.1.1 CONSTITUENTS AND FORCES

The standard model (SM) of particle physics consists of elementary particles that make up our universe, carry the fundamental forces, and give mass to most things around us. Each of those elementary particles is a fundamental building block of nature, which means that two elementary particles of the same kind are indistinguishable perfect copies of each other. Six quarks, six leptons, four force-carrying particles, and one scalar boson form the SM as illustrated in fig. 3.1. All of them can be grouped into two categories: fermions and bosons. Fermions have half-odd integer spin, while bosons have integer spin. The six leptons, e, μ, τ and their corresponding neutrinos ν_e, ν_μ, ν_τ , and the six quarks are fermions with spin $\frac{1}{2}$. The six leptons are classified by their charge ($Q = -1$ for e^-, μ^-, τ^- and $Q = 0$ for ν_e, ν_μ, ν_τ) and lepton number, and the signs are reversed for the corresponding six anti-leptons. Therefore, counting the antiparticles as well, the SM contains 12 leptons. Up-type quarks (up (u), charm (c), top (t)) have an electrical charge of $Q = 2/3$, while the down-type quarks (down (d), strange (s), bottom (b)) have an electrical charge of $Q = -1/3$. Again, the signs are reversed for the corresponding antiparticles. Quarks are also classified by their colour charge: red, green, or blue. Hence, the SM contains 36 quarks if we count the particles, antiparticles, and different colours. When different quarks are combined, they form particles that are not elementary, known as hadrons. All isolated hadrons must be colour-neutral, which is why isolated quarks have never been observed. An equal number of quarks and antiquarks form mesons, and an odd number of quarks (antiquarks) form so-called baryons. The photon γ mediates the electromagnetic force, the bosons W^+ and W^- mediate the weak force, Z mediates the electroweak force and gluons g mediate the strong force. All elementary particles couple via the weak force, charged leptons and quarks couple also electromagnetically, and only quarks interact strongly due to their colour charge. Each gluon carries a combination of colour and anticolour. In the $SU(3)$ representation of quantum chromodynamics (the theory of strong interactions in the SM), this results in eight independent colour states. Together with the Higgs boson H , the total elementary particle content of the SM is 61 [20].

3.1.2 INPUT PARAMETER

In an ideal scenario, the ultimate theory that describes our observations would be one equation with a few input parameters, and everything would follow. The SM is far away from being that theory. Instead, the SM




Figure 3.1: Constituents of the standard model of particle physics. Image taken from ref. [19].

is a collection of renormalisable and self-consistent theories combined in the unitary product group $SU(3) \times SU(2) \times U(1)$. The quark sector is represented by quantum chromodynamics in the $SU(3)$ group, and $SU(2) \times U(1)$ is the representation of the unification of the weak force and the electromagnetic force in the Weinberg, Salam, and Glashow theory [21–23]. The fundamental objects are quantum fields, and particles are considered to be the excited states of those quantum fields. The ingredients that make the SM are the Dirac equation that describes the dynamics of fermions, quantum field theory [24], the principle of local gauge invariance, the Higgs mechanism [25, 26], and many experimental results. It is necessary to input a total of 26 parameters [27] manually.

The input parameters of the SM are:

- The fermion masses m_f or their Yukawa couplings to the Higgs field y_f , as they are related via the Higgs vacuum expectation value v : $m_f = y_f v / \sqrt{2}$. In total, there are twelve fermions to parameterise: $m_{\nu_1}, m_{\nu_2}, m_{\nu_3}, m_e, m_\mu, m_\tau, m_u, m_d, m_c, m_s, m_t, m_b$. $m_{\nu_1}, m_{\nu_2}, m_{\nu_3}$ are the three neutrino masses that correspond to the neutrino mass states ν_1, ν_2, ν_3 .
- The three coupling constants g' ($U(1)$ gauge coupling, electromagnetic), g_w ($SU(2)$ gauge coupling, weak), and g_s ($SU(3)$ gauge coupling, strong).
- The parameters that describe the Higgs potential: The Higgs mass m_H and the Higgs vacuum expectation value v .
- The parameters of the Cabibbo–Kobayashi–Maskawa (CKM) matrix [28, 29] in the Wolfenstein parameterisation [30] λ, A, ρ, η that connects the quark mass and weak-interaction eigenstates, and the

mixing angles and complex phase of the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [31, 32] $\Theta_{12}, \Theta_{13}, \Theta_{23}, \delta$ that connects the neutrino mass and flavour eigenstates.

- A CP-violating phase in the QCD Lagrangian Θ_{CP} that is known to be close to zero [27].

3.1.3 SYMMETRIES

The SM is a chiral theory where fermions ψ transform according to their left-handed ψ_L or right-handed ψ_R representation (γ^5 is the fifth gamma matrix defined as $\gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3$, with $\gamma^0, \gamma^1, \gamma^2, \gamma^3$ being the gamma matrices or Dirac matrices):

$$\psi_L = \frac{1}{2}(1 - \gamma^5)\psi \quad \psi_R = \frac{1}{2}(1 + \gamma^5)\psi \quad (3.1)$$

Left-handed particles are grouped into weak-isospin (T) doublets in $SU(2)_L$ (L for left), and right-handed particles form singlets in $U(1)_Y$ (with the weak hypercharge Y). Neutrinos are the only leptons that do not have right-handed counterparts in the SM.

The SM has many symmetries that are connected to conservation laws according to Noether’s theorem [33]. The beauty of Noether’s theorem is that each conservation law follows directly from a given symmetry of a physical system. For example, the conservation of energy follows from the uniformity of time. Global symmetries in the SM are energy, momentum, angular momentum and baryon and lepton number conservation. The local gauge symmetry in the $SU(3)_C \times SU(2)_L \times U(1)_Y$ group leads to the conservation of the colour charge C , the weak isospin T and its third component T_3 , the electric charge Q , and the weak hypercharge Y . Those quantities are intertwined by the Gell-Mann–Nishijima formula: $Q = T_3 + \frac{1}{2}Y$ [34–36].

3.1.4 MASS

The Higgs mechanism explains how the gauge bosons $W^\pm$ and Z , as well as the fermions, acquire their mass by spontaneously breaking the $SU(2)_L \times U(1)_Y$ symmetry. This happens when the Higgs potential $V(\phi) = \mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2$ reaches one of its infinite minima satisfying $\phi^\dagger\phi = v^2/2 = -\mu^2/(2\lambda)$ for $\mu^2 < 0$ [27]. λ is the Higgs self-coupling and ϕ the Higgs doublet. After spontaneous symmetry breaking, the initial Yukawa coupling terms with the Higgs-boson split into two parts: one actual interaction term and one mass term. The masses of the gauge bosons follow directly from the Higgs vacuum potential and the coupling constants:

$$m_W = \frac{1}{2}g_W v \quad m_\gamma = 0 \quad m_Z = \frac{1}{2}v\sqrt{g_W^2 + g'^2} \quad (3.2)$$

As mentioned earlier, the Higgs vacuum potential connects the fermion Yukawa couplings and masses. The Higgs-boson mass is related to the free parameters of the Higgs potential $m_H^2 = 2\lambda v^2$. Figure 3.2 shows the fermion masses for the three generations. The origin of the wide range of masses and the pattern within the different fermion generations is unknown. The fermion mass scale differences are related to the hierarchy




Figure 3.2: Summary of the fermion masses split by generation. The neutrino masses are shown approximately assuming the normal hierarchy $m_1 < m_2 < m_3$. Image taken from ref. [27].

problem in physics: The question of why the Planck energy scale $\mathcal{O}(10^{19})$ GeV, at which quantum effects of gravity become significant, and the electroweak energy scale $\mathcal{O}(10^2)$ GeV are so different [37].

3.1.5 SUCCESSES AND LIMITATIONS

Despite its limitations, the successes of the SM are undeniable. Figure 3.3 summarises the total cross-section measurements of different particle production processes in pp collisions by ATLAS [38]. The measured cross-sections agree very well with the SM theory prediction shown in grey over a wide range of processes.

However, the SM will not be the end of the story. Gravity is not included in the SM. We do not know what dark matter is made of, and the SM does not contain a dark matter candidate that would fit the cosmological constraints, assuming dark matter is a particle. Furthermore, the CP-violating terms in the SM are not big enough to explain the baryon asymmetry in the universe. The SM has too many parameters that must be given by hand: Why are there three generations of quarks and leptons, and why is their mass scale so different? Why are the neutrino masses so much smaller than the other lepton masses? And finally, is the Higgs boson an elementary particle [39]? The observation of neutrino oscillations causes the most pressing need for physics beyond the SM. In the upcoming section, we will delve into that topic.




Figure 3.3: Summary of SM cross-section measurements performed by ATLAS. Image taken from ref. [38].

3.2 HEAVY NEUTRAL LEPTON MODELS

Let me start with a disclaimer: the landscape of models that involve right-handed neutrinos is vast. Most can be categorised and derived from a general addition to the SM Lagrangian. I will outline the big picture of heavy neutral lepton (HNL) models in this section, but I won't be able to dive deep into the various models. I received most of the inspiration and references for this section from a recent review about the 'present and future of heavy neutral leptons' [6], which I highly encourage you to read.

Once the broad description is done, I'll converge towards the realistic physics model used in this thesis to interpret the results, the Quasi-Dirac model [40]. But what actually motivates the search for HNLs?

3.2.1 OVERVIEW

At the time of building the SM, neutrinos have been assumed to be massless. Therefore, only left-handed neutrinos have been added to the SM, as there was no need for additional right-handed counterparts that do not couple. If those right-handed counterparts had been added, neutrino mass terms could have been implemented via Dirac mass terms analogous to the charged fermions. There is only one direct evidence for beyond-the-standard-model physics (BSM) measured in the laboratory to date: the observation that neutrinos are not massless. The earliest evidence for neutrino masses was the observed deficit in solar neutrinos [3] in 1968 at the Brookhaven National Laboratory. In 1998, the Super-Kamiokande Collaboration published the first evidence for the oscillation of atmospheric neutrinos [4]. This was followed by the publication of the direct evidence that neutrinos change their flavour by the SNO Collaboration in 2002 [5].

Meanwhile, we have reasonably accurate measurements of the neutrino mass splittings Δm_{21}^2 , $|\Delta m_{31}^2|$, $|\Delta m_{32}^2|$ ($\Delta m_{ij}^2 = m_{\nu_i}^2 - m_{\nu_j}^2$) and mixing angles Θ_{23} , Θ_{12} , Θ_{13} . The unknown sign of the mass splittings $|\Delta m_{31}^2|$, $|\Delta m_{32}^2|$ determines the neutrino mass hierarchy:

$$\begin{aligned} \text{normal hierarchy (NH): } m_1 < m_2 < m_3 \quad \Delta m_{31}^2 > 0 \\ \text{inverted hierarchy (IH): } m_3 < m_1 < m_2 \quad \Delta m_{32}^2 < 0 \end{aligned} \tag{3.3}$$

Δm_{21}^2 is known very precisely from neutrino oscillations, while the sign of $|\Delta m_{31}^2|$ is unknown, and its value is less precisely measured than Δm_{21}^2 . Two non-zero neutrino mass splittings imply that at least two of the three SM neutrinos have non-zero mass. Knowing the mass splittings means that we have only one unknown parameter left, the minimal neutrino mass $m_{\min}$:

$$\begin{aligned} \text{NH: } m_1 &= m_{\min} \quad m_2 = \sqrt{m_{\min}^2 + \Delta m_{21}^2} \\ &\quad m_3 = \sqrt{m_{\min}^2 + \Delta m_{31}^2} \\ \text{IH: } m_3 &= m_{\min} \quad m_1 = \sqrt{m_{\min}^2 + |\Delta m_{32}^2| - \Delta m_{21}^2} \\ &\quad m_2 = \sqrt{m_{\min}^2 + |\Delta m_{32}^2|} \end{aligned} \tag{3.4}$$

Using the relations in eq. (3.4) and the measured values of the mass splittings (global fit values from [41]: $\Delta m_{21}^2 = 7.42_{-0.20}^{+0.21} \cdot 10^{-5} \text{ eV}^2$, $\Delta m_{31}^2 = 2.517_{-0.028}^{+0.026} \cdot 10^{-3} \text{ eV}^2$ (NH), and $\Delta m_{32}^2 = -2.498_{-0.028}^{+0.028} \cdot 10^{-3} \text{ eV}^2$ (IH)), we can constrain the mass of all neutrinos to be above $\sum m_j \geq 60 \text{ meV}$ for NH and $\sum m_j \geq 100 \text{ meV}$ for IH [6].

Other constraints on the neutrino mass scale come from experiments such as KATRIN, which places an upper limit on $m_\nu = \sqrt{\sum_{j=1,2,3} |U_{ej}|^2 m_j^2} < 0.8 \text{ eV}$ [42], where U_{ej} are the electron mixing elements of the PMNS matrix. Constraints from cosmological searches set upper bounds on the sum of all neutrino masses $\sum m_j < 0.12 \text{ eV}$ [43, 44].

Now, given the constraints, we know the ballpark of the neutrino mass scale, and we know for sure that neutrinos are not massless. Still, the smallness of neutrino masses raises important questions [6]: What is the nature of neutrinos? Are they Dirac or Majorana particles? How light are they exactly? What is the reason for their small masses? Is there a CP violation in the neutrino sector, and if so, what is the value of the CP violating phase δ in the PMNS matrix? What are the values of the mixing angles Θ_{ij} , and is the standard three neutrino picture correct?

A possible explanation for the smallness neutrino masses is the Seesaw mechanism [45–52] that postulates the existence of some electroweak-singlet neutrino interaction states. But let's get back to that later.

The general goal is to extend the SM to a full theory of particles and interactions. In the SM, neutrinos are left-handed Dirac particles, but since they are neutral, they could be both Dirac or Majorana, dependent on the requirement of lepton number conservation. If they were Majorana particles, they would fulfil the Majorana condition: $\nu = \nu^C \equiv C \bar{\nu}^T$, where C is the charge conjugation matrix. Neutrinos and antineutrinos would be indistinguishable.

So, how do we get the neutrino masses in the SM? One could add a Dirac mass term to the SM Lagrangian, analogous to the charged fermions in the SM, that would add right-handed counterparts ν_R to the left-handed SM neutrinos ν_L :

$$\text{Dirac mass term } \bar{\nu} m_D \nu = \bar{\nu}_L m_D \nu_R + h.c. \quad (3.5)$$

The Dirac term would conserve the lepton number. A second option would be the addition of a Majorana mass term:

$$\text{Majorana mass term } \frac{1}{2} \bar{\nu}^C m_M \nu = -\frac{1}{2} \nu_L^T C^\dagger m_M \nu_L + h.c. \quad (3.6)$$

Both mass terms are forbidden in the SM because right-handed neutrinos do not exist in the SM, Majorana terms that use only left-handed SM neutrinos violate the SM gauge invariance, and, finally, the SM in its minimal form does not allow neutrino masses [6].

There are many proposed models to unravel the situation. The simplest one would be to add right-handed sterile neutrinos ν_R as SM gauge singlets to the SM since the SM gauge group permits Yukawa interactions between ν_R , the lepton doublet $L = (\nu_L^T, l^T)^T$, and the Higgs doublet H :

$$-\mathcal{L}_y = \bar{L} F \cdot \tilde{H} \nu_R + h.c. \quad (3.7)$$

where $\tilde{H} = i\sigma_2 H^*$ and F is the matrix of Yukawa couplings. After electroweak symmetry breaking $\langle \tilde{H} \rangle = (v/\sqrt{2}, 0)^T$, with the Higgs vacuum expectation value v :

$$-\mathcal{L}_y \xrightarrow{\langle \tilde{H} \rangle \neq 0} -\mathcal{L}_{\text{Dirac}} = \frac{v}{\sqrt{2}} \bar{\nu}_L y_\nu \nu_R + h.c. \quad (3.8)$$

with the Yukawa interaction term y_ν . Also, here, the lepton number is conserved. The disadvantage is that the coupling y_ν needs to be very small compared to the other fermion Yukawa couplings to reach the observed neutrino mass scale.

Let's go one step further. The SM stays gauge invariant when adding the Weinberg operator [53], which has the mass dimension five and is a low-energy effective term with the scale Λ and coupling $c^{[5]}$:

$$\mathcal{L}_{\text{M,BSM}} = \frac{c^{[5]}}{\Lambda} L^T \cdot \tilde{H}^* C^\dagger \tilde{H} \cdot L + h.c. \quad (3.9)$$

The Weinberg operator leads to Majorana masses $m_\nu = -v^2 c^{[5]}/\Lambda$ after electroweak symmetry breaking as it violates the lepton number symmetry:

$$\mathcal{L}_{\text{M,BSM}} \xrightarrow{\langle \tilde{H} \rangle \neq 0} \frac{c^{[5]} v^2}{2\Lambda} \nu_L^T C^\dagger \nu_L + h.c. \quad (3.10)$$

New particles arise with masses $M \approx \Lambda$. At the same time, the conservation of Lorentz and gauge symmetries leaves three options for the nature of these particles and the origin of the Weinberg operator in eq. (3.9) [6]: Seesaw type I mechanism with singlet fermions [45, 46, 48, 49, 51], Seesaw type II mechanism with heavy triplet scalars [51, 54], and Seesaw type III mechanism with triplet fermions [55, 56]. Figure 3.4 summarises the diagrams that contribute at tree level to the neutrino masses in the three Seesaw scenarios.

We call the new fermionic states heavy neutral leptons (HNLs). The scale of those HNLs ranges from a few MeV to very high energies. The scale in eq. (3.10) can be fine-tuned to accommodate the smallness of the neutrino masses.

The neutrino masses are suppressed by v/Λ in the Weinberg operator, so if we want to keep the couplings $c^{[5]}$ at an accessible level ($\mathcal{O}(1)$ is considered natural), we can set the scale Λ to very large values ($\sim 10^{14}$ GeV) close to the GUT scale. Even leptogenesis can be realised in this formulation. The downside here is that because of its large scale, the model cannot be tested because the new particles would be too heavy. Another




Figure 3.4: Diagrams contributing to neutrino masses in the three Seesaw scenarios at tree level. $\langle H \rangle$ indicates the vacuum expectation value of the neutral component of the Higgs field. N_R is a singlet fermion (with mass M_R), Δ is a scalar triplet and Σ is the neutral component of a fermion triplet, with mass M_Σ . Image and caption taken from [6].

thing to note is that the Higgs mass corrections become too large without a stabilising mechanism for the electroweak symmetry-breaking scale (such as Supersymmetry) [57–60].

Another approach would be to adjust the couplings $c^{[5]}$ to very small values ($c^{[5]} \ll 1$). This way, the scale Λ can be lower, and HNLs are within the reach of current experiments if the coupling is not too low. In this scale (HNL masses between $\mathcal{O}(\text{eV})$ and $\mathcal{O}(\text{TeV})$), the observed baryon asymmetry in the universe can be explained via the ARS mechanism (Baryogenesis via neutrino oscillations) [61, 62] and resonant leptogenesis [63–65].

The third option would be to utilise quasi-conserved symmetries. The Majorana neutrino masses are zero if the lepton number symmetry is preserved. A small breaking of such a lepton number symmetry would generate small Majorana masses. The small lepton number symmetry breaking is used in the inverse [66, 67], linear [68, 69], and extended Seesaw mechanisms. Here, the lepton number violation is small, and the Majorana HNLs have a very low mass difference, which makes them quasi-Dirac. The couplings here can be larger, while the HNLs are within the MeV-TeV scale. Therefore, collider searches can probe this HNL parameterisation.

3.2.2 TYPE I SEESAW MECHANISM

Now, let's take one step back and write down the type I Seesaw [45, 46, 48, 49, 51] Lagrangian for the addition of $I, J = 1, 2, 3$ right-handed HNLs, denoted by N , to the SM Lagrangian $\mathcal{L}_{\text{SM}}$:

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + i\bar{N}_I \gamma^\mu \partial_\mu N_I - F_{\alpha I} \bar{L}_\alpha \tilde{H} N_I - \frac{M_{IJ}}{2} \bar{N}_I^C N_J + h.c. \quad (3.11)$$

$F_{\alpha I}$ is the matrix of Yukawa couplings with the flavour index $\alpha = e, \mu, \tau$. The SM neutrino masses scale as $m_\nu \sim \frac{F^2 v^2}{M}$ after electroweak symmetry breaking. Assuming $F \sim \mathcal{O}(1)$ easily sends the mass scale $M \rightarrow 10^{15} \text{ GeV}$ to the GUT scale to get the correct SM neutrino mass scale (as discussed before). An additional mass suppression can be acquired from an approximate symmetry [70, 71] analogous to the inverse and linear Seesaw

mechanisms. The lepton number is violated by the mass term of the right-handed Majorana neutrino, which could be responsible for the observed matter-antimatter asymmetry in the universe [72]. The connection between neutrino masses m_ν , Majorana masses M , mixing angles Θ , and the PMNS matrix U_{PMNS} is given by the Seesaw formula:

$$m_\nu = \frac{v^2}{2} F M^{-1} F^T = \Theta M \Theta^T = U_{\text{PMNS}} M U_{\text{PMNS}}^T \quad (3.12)$$

The mixing angles between the active neutrinos and the HNLs that connect the HNLs to the visible SM sector then follow from:

$$\Theta = \frac{vF}{\sqrt{2}} M^{-1} \quad (3.13)$$

The Seesaw hypothesis $Fv \ll M$ implies that the mixing angles are too small for HNL detection with current experiments. However, we can apply the symmetry considerations again and look at scenarios where the lepton number is approximately conserved to retain larger mixing angles [67, 70, 71]. The SM neutrino masses then follow a slightly modified Seesaw equation with small lepton number breaking parameters μ_2 and ϵ [6]:

$$m_\nu = v^2 \epsilon (Y' M^{-1} Y^T + Y M^{-1} Y'^T) - v^2 Y \mu_2 M^{-1} Y^T \quad (3.14)$$

Y, Y' are Yukawa coupling matrices and M is the HNL mass matrix with the two cases $\mu_2 \neq 0, \epsilon = 0$ for the inverse Seesaw mechanism and $\epsilon \neq 0, \mu_2 = 0$ for the linear Seesaw mechanism. The lepton number breaking parameters μ_2 and ϵ lead to light neutrino masses together with large enough couplings Y to get HNLs back into reach of experiments. Moreover, the mixing angles are much larger than in the standard Seesaw scenario.

If we want to get ahead of ourselves and speak of HNL discovery, we would want to measure all HNL parameters to test the minimal model fully [73]. Here, a parameterisation can be used that is related to the Seesaw formula for the mixing Θ between the SM neutrinos and the HNLs:

$$\Theta \sim i U_{\text{PMNS}} m^{1/2} R^\dagger M^{-1/2} \quad (3.15)$$

where U_{PMNS} is the PMNS matrix, m the light neutrino mass matrix, R a $3 \times n$ orthogonal matrix related to the HNL sector (n HNL states), and M is the HNL mass matrix.

The, by far, most quoted model in the HNL community is the neutrino minimal standard model (νMSM) [7, 8] as shown schematically in fig. 3.5. Here, the number of HNLs added to the SM is symmetrised to the number of SM neutrinos. So we add exactly three HNLs $N_{1,2,3}$ to the SM (using the Lagrangian as shown in eq. (3.11)). Two of them, $N_{2,3}$, are heavy ($\sim \text{GeV}$) and responsible for generating the masses of the SM neutrinos via the type I Seesaw mechanism. They can also generate the baryon asymmetry in the universe, but only if their masses are close to each other [7, 61]. The leftover HNL N_1 is light and serves as a bystander to be the candidate for warm dark matter. N_1 is too light to be detected by a collider experiment.




Figure 3.5: Schematic illustration of the particle content of the neutrino minimal SM (ν MSM) [7, 8] split up in their left-handed and right-handed components. The SM is extended by three right-handed heavy neutral leptons $N_{1,2,3}$.

The ballpark for the mass of N_1 is the keV range, but some references state 100 eV [74], 0.1-10 keV [75], 1 keV-10 MeV [76], and < 3.5 keV [77]. Another reason why N_1 is out of range for experiments is its Yukawa coupling, which needs to be tiny ($5 \cdot 10^{-15} - 4 \cdot 10^{-13}$) [78]. Overall, the ν MSM is a consistent theory to very high scales, potentially up to the Planck scale $\mathcal{O}(10^{19})$ GeV [9].

To wrap up the journey through the lands of HNL models, I want to mention beyond-minimal models briefly. HNLs don't need to be added as a minimal extension to the SM. They can also be a part of a potential hidden sector with new particles and interactions. Examples here are dark sector models, where a hidden $U(1)_x$ symmetry extends the SM gauge group [79]. Another example are the left-right symmetric models that attempt to restore the electroweak parity violation via a gauge group extension $SU(2)_L \times SU(2)_R$ at high energies [80, 81]. Originally, the Seesaw mechanism, together with right-handed neutrinos, arose from the context of larger gauge groups. Examples here are grand unification theories of $SU(5)$ [47], $SO(10)$ [48], as well as parity restoration theories [45] and partial unification theories as $SU(4)_C \times SU(2)_L \times SU(2)_R$ [82].

The landscape of HNL models is rich, and we profit from the constant curiosity of the theory community. I omitted more exotic models here to keep the tale of models short and the storyline as straight as possible.

3.2.3 A QUASI-DIRAC PAIR OF HEAVY NEUTRAL LEPTONS

Now, let us converge on the Quasi-Dirac model, which is used in this work and falls into the third category of quasi-conserved symmetries, as discussed in the previous section. In the Quasi-Dirac model, we consider a

pair of right-handed Majorana neutrinos (still HNLs) with highly degenerate mass $M_1 \approx M_2$, which makes both HNLs seem like a single (Quasi-)Dirac particle:

$$\Delta M \equiv |M_1 - M_2| \ll M_N = \frac{M_1 + M_2}{2} \quad (3.16)$$

Hence, the lepton number is a quasi-conserved symmetry since the HNL pair is a Quasi-Dirac particle that violates the lepton number conservation only slightly. As in the previous type I Seesaw models, we start with the Lagrangian in eq. (3.11). After electroweak symmetry breaking ($\langle \tilde{H} \rangle = (v/\sqrt{2}, 0)^T$), the Yukawa interaction term $-F_{\alpha I} \bar{L}_\alpha \tilde{H} N_I$ becomes a Dirac mass term $[M_D]_{\alpha I} = F_{\alpha I} v$ that couples the left and the right-handed components of the neutrinos. The right-handed neutrinos have an additional Majorana mass term $M_{IJ} = [M_M]_{IJ} = M_I \delta_{IJ}$. We apply again the Seesaw hypothesis, this time with the additional requirement to be below the electroweak energy scale $\Lambda_{EW} = \mathcal{O}(10^2)$ GeV [83]:

$$|F_{\alpha I}|v \ll |M_I| \leq \Lambda_{EW} \quad (3.17)$$

The full mass matrix:

$$\hat{M} = \begin{pmatrix} 0 & M_D \\ M_D^T & M_M \end{pmatrix} \quad (3.18)$$

can be diagonalised $\hat{U}^\dagger \hat{M} \hat{U}^* = \hat{M}^{\text{diag}}$ with the unitary matrix $\hat{U}$ ($\hat{U}^\dagger \hat{U} = \hat{U} \hat{U}^\dagger = 1$):

$$\hat{U} = \begin{pmatrix} U_{\text{PMNS}} & \Theta \\ -\Theta^\dagger U_{\text{PMNS}} & 1 \end{pmatrix} \quad (3.19)$$

with the matrix of mixing angles Θ and the PMNS matrix U_{PMNS} . The active neutrino mass eigenstates follow from

$$\text{diag}(m_1, m_2, m_3) = U_{\text{PMNS}}^\dagger M_\nu U_{\text{PMNS}}^* \quad (3.20)$$

with the Seesaw mass matrix $M_\nu = -M_D M_M^{-1} M_D^T$. Together with the three light neutrino mass eigenstates ν_i ($i = 1, 2, 3$) we get two heavy eigenstates $\mathcal{N}_I$ ($I = 1, 2$) with masses M_1 and M_2 . Those heavy mass eigenstates are the additional right-handed Majorana neutrinos or HNLs.

One thing to note is that there are no direct arguments that HNLs must be Majorana particles. They could as well be Dirac particles. However, the urge to keep the extension of the SM renormalisable automatically leads to the Majorana mass term, as discussed earlier. Therefore, HNLs are usually assumed to be Majorana particles.

In the particular parameterisation of two additional HNLs, the lightest neutrino is assumed to be massless ($m_{\min} = 0$) [84].

The flavour eigenstates of the active neutrinos $\nu_{L\alpha}$ ($\alpha = e, \mu, \tau$) can be written as a linear combination of the five mass eigenstates ν_i and $\mathcal{N}_I$ [40, 83]:

$$\nu_{L\alpha} = U_{\alpha i}^{\text{PMNS}} \nu_i + \Theta_{\alpha I} \mathcal{N}_I^C \quad (3.21)$$

where the mixing angles $\Theta_{\alpha I} \approx v F_{\alpha I} / M_I$. The Yukawa couplings $F_{\alpha I}$ can be written in the Casas-Ibarra parameterisation [85, 86] with no loss of generality [83]:

$$F = \frac{i}{v} V D_\nu^{1/2} R D_\mathcal{N}^{1/2} \quad (3.22)$$

with $D_\nu = \text{diag}(m_1, m_2, m_3)$, $D_\mathcal{N} = \text{diag}(M_1, M_2)$, $m_1 = 0$ for the normal neutrino mass hierarchy, $m_3 = 0$ for the inverted neutrino mass hierarchy, and the orthogonal 3×2 matrix R ($R^T R = 1$) that depends on the active neutrino mass hierarchy:

$$R_{\text{NH}} = \begin{pmatrix} 0 & 0 \\ \cos \omega & \sin \omega \\ -\sin \omega & \cos \omega \end{pmatrix} \quad R_{\text{IH}} = \begin{pmatrix} \cos \omega & \sin \omega \\ -\sin \omega & \cos \omega \\ 0 & 0 \end{pmatrix} \quad (3.23)$$

with an arbitrary complex number ω . The matrix V is the PMNS matrix U_{PMNS} multiplied by an additional Majorana phase η :

$$V = U_{\text{PMNS}} \times \text{diag}(1, e^{i\eta}, 1) \quad (3.24)$$

where $\eta = (\alpha_{21} - \alpha_{31})/2$ (NH) or $\eta = \alpha_{21}/2$ (IH), and $\eta \in [0, 2\pi[$. The PMNS matrix is written in terms of the three active neutrino mixing angles Θ_{12} , Θ_{23} , Θ_{13} and a Dirac phase δ :

$$\begin{aligned} U_{\text{PMNS}} &= \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \\ &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \end{aligned} \quad (3.25)$$

with $s_{ij} = \sin \Theta_{ij}$ and $c_{ij} = \cos \Theta_{ij}$.

So, in summary, we add eleven new parameters [40] to the SM Lagrangian in eq. (3.11): m_1, m_2, m_3 , Θ_{12} , Θ_{23} , Θ_{13} , and δ for the active neutrino sector and M_1, M_2, η , and ω for the HNL sector. As we learned in eq. (3.4), m_1, m_2, m_3 can also be expressed in terms of Δm_{21}^2 , Δm_{31}^2 (or $|\Delta m_{32}^2|$), and $m_{\min}$.

Now, actual observable quantities are the coupling strength of the HNL to a particular neutrino flavour $|U_\alpha|^2$ and the total coupling $|U_{\text{tot}}|^2$:

$$\begin{aligned} |U_\alpha|^2 &\equiv \sum_I |\Theta_{\alpha I}|^2 \quad |U_{\text{tot}}|^2 \equiv \sum_{\alpha, I} |\Theta_{\alpha I}|^2 \\ |U_e|^2/|U_{\text{tot}}|^2 + |U_\mu|^2/|U_{\text{tot}}|^2 + |U_\tau|^2/|U_{\text{tot}}|^2 &= 1 \end{aligned} \quad (3.26)$$

The total coupling can be expressed as (expressions for $|U_\alpha|^2$ are more complicated and given in ref. [87]):

$$|U_{\text{tot}}|^2 = \frac{\sum_i m_i}{M_N} \cosh(2\Im\omega) \quad M_N = \frac{M_1 + M_2}{2} \quad (3.27)$$

The observed data related to the active neutrino sector [41] constrains the allowed regions in the $(|U_e|^2, |U_\mu|^2, |U_\tau|^2)$ coupling space, as shown in fig. 5.7. Low-scale leptogenesis also motivates the choice of $\Delta M \ll M_N$ [88]. The theory has an approximate global $U(1)$ lepton number symmetry, which is exact in the $M_1 = M_2$ and $\Theta_{\alpha 1} = \pm i\Theta_{\alpha 2}$ limit where both HNLs form a single Dirac particle $\psi = N_1 + N_2^C$ and the active neutrinos are massless. HNL oscillations, which can also occur in the discussed model, are further quantified in ref. [40].

Two very relevant observable HNL parameters are the total coupling $|U_{\text{tot}}|^2$ and the average HNL mass M_N . Figure 3.6 shows the allowed range of those parameters consistent with both the Seesaw mechanism and leptogenesis for the normal and inverted neutrino mass hierarchies assuming the presence of two HNLs¹ as in the Quasi-Dirac model. The lower constraint on the total coupling is the type I Seesaw bound corresponding to the requirement to reproduce the light neutrino masses. The upper constraint stems from leptogenesis (in particular, baryogenesis via oscillations and resonant leptogenesis; see ref. [64] for details). A green colour bar in fig. 3.6 indicates the largest allowed value of the mass splitting normalised to the average mass $\Delta M/M_N$ in the allowed regions. At the edges of the allowed parameter space, the normalised mass splitting $\Delta M/M_N$ has to be smaller than 10^{-6} . The mass splitting ΔM sets the scale for the oscillations that violate CP and that lead to a lepton asymmetry [64], contributing to the observed matter-antimatter asymmetry in the universe.

This work exploits the Dirac-limit ($\Delta M = 0$) and the Majorana-limit ($\Delta M \neq 0$) of the HNL pair. Investigating both cases might shed some light on the unknown parameter ΔM and its impact on the model. The application of the Quasi-Dirac model to the data analysis performed in this work will be discussed in detail in section 5.3.

¹Similar constraints that allow a wider parameter space exist for the case of three HNLs [65], where again very similar masses for all three HNLs are assumed. The case of three HNLs is also interesting, as the lightest SM neutrino does not have to be massless ($m_{\text{min}} \geq 0$). A summary of all theoretical constraints, experimental constraints, and expected limits from future experiments will be shown at the end of this work in fig. 8.2.




Figure 3.6: The range of the total coupling $U^2 = |U_{\text{tot}}|^2$ consistent with both the Seesaw mechanism and leptogenesis as a function of the average HNL mass $M = M_N$ for the case of two HNLs. The black solid lines show the results obtained with the full kinetic equations and vanishing initial conditions for HNLs. The blue dashed lines correspond to thermal initial conditions. In this regime, the freeze-in does not contribute to the asymmetry generation. The red dotted line neglects the effect of the expansion of the universe on the distribution of the heavy neutrinos. In this case, the freeze-out cannot contribute, and the asymmetry is generated during freeze-in. The colour contours represent the largest allowed value of the mass splitting $\Delta M/M_N$. ΔM and M_N are defined in eq. (3.16). Within the white regions, the mass splitting has to be smaller than 10^{-6} . The left (right) panel shows the case of the normal (inverted) neutrino mass hierarchy. Figures and modified caption are taken from ref. [64].

4 THE EXPERIMENT

4.1 LARGE HADRON COLLIDER

The Large Hadron Collider (LHC) [90, 91] is the most powerful tool we have to date to explore the world of particle physics. The 26.7 km long circular tunnel is located 45-170 m below the surface, crossing the state borders of Switzerland and France close to Geneva with a beautiful view of Mont Blanc. Four caverns that are connected to the surface host the four main experiments ATLAS [92], CMS [93], LHCb [94], and ALICE [95] at the interaction points where the two counterrotating proton beams cross, as shown in fig. 4.1. The total cost of the LHC, including detectors, detector R&D, and pre-operation, is 4332 MCHF [96].

The tunnel that was originally constructed between 1984 and 1989 for the large electron-positron collider (LEP) consists of eight straight sections and eight arches. In 2000, LEP was closed to make space for the LHC. The challenge to reuse the LEP tunnel led to a two-in-one superconducting magnet to house both counter-rotating proton beams in the same tunnel. Those dipole magnets house both proton beams instead of having a separate magnet for each beam. To reach the required kinetic beam energy of $E = (\gamma - 1)mc^2 = 7$ TeV, with $\gamma = (1 - \beta^2)^{-1/2}$, that leads to the desired centre-of-mass energy $\sqrt{s} = 14$ TeV, the protons need to be accelerated to $\beta = v/c = 0.999999991$ times the speed of light c . The total energy E_i and momentum $\mathbf{p}_i$ of the initial state particles i define the Lorentz invariant centre-of-mass energy:

$$\sqrt{s} = \sqrt{\left(\sum_i E_i\right)^2 - \left(\sum_i \mathbf{p}_i\right)^2} \quad (4.1)$$

The design beam energy of 7 TeV requires a magnetic field strength of 8.33 T to hold the protons in orbit, generated by magnets using NbTi Rutherford cables cooled with superfluid helium down to 1.9 K [91]. The protons are injected in bunches, where each bunch contains about $1.1 \cdot 10^{11}$ protons. Once filled, the LHC hosts 2808 proton bunches with an average spacing of 25 ns. Two independent radio-frequency (RF) cavity systems accelerate and shape the two counterrotating LHC beams into their bunch structure at an operating frequency of 400.8 MHz [91].

The target luminosity of the LHC is $L = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and depends on several parameters [91]:

$$L = \frac{N_b^2 n_b f_r \gamma}{4\pi \varepsilon_n \beta^*} F \quad (4.2)$$




Figure 4.1: Overview of the LHC and its underground location in Switzerland and France. Image taken from ref. [89].

Those parameters are the number of protons per bunch N_b , the number of bunches per beam n_b , the revolution frequency f_r , the relativistic factor γ , the normalised transverse beam emittance ε_n , the beta function at the collision point β^* , and a geometric reduction factor F that depends on the crossing angle of the two beams at the interaction point.

Figure 4.2 shows the CERN accelerator complex. The protons are accelerated in different stages. First, LINAC4¹ [99] reaches a beam energy of 160 MeV, and then, the BOOSTER [100] goes up to 1.4 GeV, followed by the PS synchrotron [101] to reach 25 GeV. The SPS synchrotron [102] reaches 450 GeV and marks the final injection stage before the LHC. Two 2.5 km transfer tunnels connect the LHC to the SPS accelerator. The LHC then ramps up the beam energy from 450 GeV to 7 TeV, which takes about 20 minutes with the two independent RF accelerating systems [91].

The LHC was initially designed for a beam energy of 7 TeV, but LHC Run2 was done with 6.5 TeV only [103]. The reason behind this is the magnet training. The magnets must be trained to reach the required magnetic field strength and electrical currents up to 12 kA without losing the superconducting state. A small amount of energy can raise the temperature of the coils such that the superconductivity state is lost, which is called a quench. The electrical current in the coils is increased until such a quench happens. Each time this process is repeated, the current can be raised further until, ideally, the target magnetic field strength is reached. During LHC operation, training some individual magnets for $\sqrt{s} = 14$ TeV would take too long, so it was decided to continue with 13 TeV for Run2 [103].

¹LINAC4 replaced LINAC2 during the long shutdown in 2019-2020. LINAC2 had accelerated protons up to an energy of 50 MeV.




Figure 4.2: Overview of the LHC and the CERN accelerator complex. Image taken from ref. [97].




Figure 4.3: Runs and technical shutdowns of the LHC and High Luminosity (HL)-LHC. Image taken from ref. [98].




Figure 4.4: Schematic of the ATLAS detector and the subdetectors in the current configuration. The NSW has been added during the technical shutdown between LHC Run2 and LHC Run3. Image taken from ref. [104].

The LHC and High Luminosity (HL)-LHC run schedule is summarised in fig. 4.3. Runs are always separated by technical shutdowns (long shutdowns (LS)) used for maintenance and detector upgrades. The main part of this work analyses data that has been recorded during Run2. When writing this thesis, we are in the middle of Run3, with a centre-of-mass energy of 13.6 TeV. After the end of Run3, a three-year technical shutdown will be used to upgrade the machine to the HL-LHC, which will most likely run until 2040 [98].

4.2 ATLAS DETECTOR

The ATLAS (A Toroidal LHC ApparatuS) detector has a cylindrical shape with a length of 46 m, a diameter of 25 m, and a total weight of 7000 tonnes. Those metrics make it the largest detector ever constructed to observe particle collisions. This section will only outline the different subdetector systems that make the ATLAS detector, ref. [92] is a very detailed description of the ATLAS detector configuration in LHC Run1, and ref. [104] describes the current ATLAS detector configuration in LHC Run3.

ATLAS uses a right-handed coordinate system with its origin at the centre of the detector and the z -axis along the beam pipe, the x -axis points to the centre of the LHC ring, and the y -axis points upward. Cylindrical coordinates are used with the radius r in the $x - y$ plane and the azimuthal angle ϕ around the beam pipe. The polar angle Θ defines the pseudorapidity $\eta = -\ln \tan(\Theta/2)$. ATLAS covers nearly all angles around the collision point.




Figure 4.5: Schematic illustration of the inner detector that consists of the pixel detector together with the insertable b-layer, the semiconductor tracker in barrel and endcap, and the transition radiation tracker in barrel and endcap. Image taken from ref. [104].

Figure 4.4 shows a schematic of the ATLAS detector and the subdetectors. Starting at the beamline, the inner detector comprises three parts: the pixel detector, the semiconductor tracker (SCT) and the transition radiation tracker (TRT). The inner detector, immersed in a $B = 2$ T axial magnetic field from the surrounding solenoid magnet, reconstructs the charged particle trajectories and measures their momentum. The electromagnetic and hadronic calorimeters start directly after the solenoid magnet. Photons and electrons deposit their energy in the electromagnetic calorimeter. In contrast, hadrons make it further out before they are absorbed in the hadronic calorimeter. The muon spectrometer is the last part of the ATLAS detector and is used to reconstruct the trajectories of muons. All subdetectors have dedicated barrel and endcap designs. The magnetic field $B = 0.5 - 1$ T required to bend the muon trajectories is created by the barrel toroid magnet and the two endcap toroid magnets. The most significant additions to the ATLAS detector between LHC Run2 and LHC Run3 are the new small wheel (NSW) and upgrades to the calorimeter electronics that improve the trigger system [104].

4.2.1 INNER DETECTOR

The inner detector, as shown in fig. 4.5 schematically, consists of the pixel detector, the semiconductor tracker (SCT), and the transition radiation tracker (TRT). All subdetectors have specific geometries for the barrel and endcap regions. Pixel and SCT form layers of concentric cylinders in the barrel and multiple discs in the




Figure 4.6: The ATLAS inner tracking detector layout for one quadrant in the $r - z$ plane, including the additional insertable B-layer (IBL) at $r = 33.5$ mm. The inner positioning tube (IPT) supports the IBL staves and separates them from the beam pipe. Next to the positions of the pixel layer, the positions of the semiconductor tracker (SCT) layers and the transition radiation tracker (TRT) are shown for both the barrel and endcap regions. The pixel and SCT detectors cover a pseudorapidity range of $|\eta| < 2.5$. Image and modified caption are taken from ref. [105].




Figure 4.7: Schematic illustration of the different calorimeter systems in ATLAS. The electromagnetic liquid argon (LAr) calorimeter splits into a barrel (EMB), an endcap (EMEC), and a forward (FCal) region. The hadronic calorimeter system consists of the Tile calorimeter in the barrel and extended barrel region and an LAr endcap (HEC) and forward (FCal) calorimeter. Image taken from ref. [104].

endcap. The TRT straw tubes are arranged parallel to the beamline in the barrel and radially in wheels in the endcap.

Figure 4.6 details the positions of the inner detector barrel and endcap layers in the $r - z$ plane. Most important for the reconstruction of displaced vertices in this work are the four layers of the pixel detector at $r = 33.5$ mm, 50.5 mm, 88.5 mm, and 122.5 mm in combination with the following four SCT layers at $r = 299$ mm, 371 mm, 443 mm, and 514 mm in the barrel region. The displaced vertex reconstruction algorithm is outlined in section 4.4.2. The insertable B-layer (IBL) is part of the pixel detector and was added just before LHC Run2 [105]. The pixel and SCT detectors cover a pseudorapidity range of $|\eta| < 2.5$. The pixel sensor has the highest granularity with a minimum pixel size of $(50 \times 400) \mu\text{m}^2$, segmented in $r - \phi$ and z , with 80.4 million readout channels. The SCT detector is made of silicon microstrips with a total of 6.3 million readout channels. The TRT is made of straw tubes that measure 4 mm in diameter and contain a $30 \mu\text{m}$ gold-plated tungsten wire in their centre. The TRT covers a pseudorapidity range of $|\eta| < 2$ with about 351000 readout channels. Combining all three subdetectors provides an outstanding pattern recognition for tracking.

4.2.2 CALORIMETERS

Figure 4.7 summarises the ATLAS calorimeter system. Dedicated calorimeter systems that vary in geometrical designs over the pseudorapidity range of $|\eta| < 4.9$ are used to detect charged and neutral particles that inter-




Figure 4.8: Schematic illustration of the different parts of the muon system in ATLAS. The new small wheel (NSW) is special since it was added recently during the long shutdown two and was not part of the ATLAS detector in LHC Run2. The muon system consists of three layers of monitored drift tube (MDT) chambers and resistive plate chambers (RPC) in the barrel region. The two endcaps comprise thin gap chambers (TGC) and MDT chambers. The barrel toroid magnet and the endcap toroid magnets create the magnetic field for the muon system. Image taken from ref. [104].

act strongly or electromagnetically. The ATLAS detector uses sampling calorimeters, where active materials follow layers of high-density materials. The basic principle is to stop the particles and force them to deposit all their energy, which is then measured by the active material. The electromagnetic calorimeter and all endcap and forward calorimeters use liquid argon (LAr) as active material. The particles ionise the argon, which is measured as an electric current. To keep the argon liquid, it needs to be cooled to -184°C .

The central hadronic calorimeter uses scintillating plastic tiles as active material and layers of steel as an absorber. The particle showers in the tiles produce photons that are transmitted through wavelength-shifting fibres to the read-out photomultipliers.

The different LAr calorimeters use different materials as absorbers: the electromagnetic LAr barrel and endcap calorimeters use lead, the electromagnetic LAr forward and hadronic LAr endcap calorimeters use copper, and the LAr hadronic forward calorimeters use tungsten. The calorimeter system is an essential input for triggering signatures left by electrons, photons, taus, jets, and missing transverse energy.

4.2.3 MUON SPECTROMETER

Figure 4.8 highlights the different constituents of the muon system. Two large superconducting air-core magnets deflect the muon tracks: the barrel toroid magnet in the pseudorapidity range of $|\eta| < 1.4$ and the two smaller endcap toroids in the pseudorapidity range of $1.6 < |\eta| < 2.7$. A superposition of magnet fields fills the transition region in the pseudorapidity of $1.4 < |\eta| < 1.6$. The $B = 0.5 - 1$ T strong magnetic field is mostly orthogonal to the muon trajectories and constantly monitored by about 1800 Hall sensors to account for perturbations. As with the other subdetector systems, the muon system splits into barrel and endcap-specific geometries.

All muon detectors are filled with gas mixtures. If a muon track crosses, the gas ionises, and electrons drift towards the anode (wire), where they induce a measurable signal. The types of muon chambers used for tracking and triggering are shown in fig. 4.8. The new small wheel (NSW) was added between LHC Run2 and LHC Run3 and consists of small strip thin gap chambers (sTGCs) and micro mesh gaseous structures (Micromegas). Both detector types are fast enough for triggering and precision tracking. All chambers must be aligned with a precision of about $30\mu\text{m}$ to reconstruct the muon trajectories. Twelve thousand alignment sensors monitor possible deviations from straight lines based on optical monitoring. Muon tracks with high transverse momentum, formed from the x and y components of the momentum vector ($p_T = \sqrt{p_x^2 + p_y^2}$), are also used to align barrel and forward regions. The resistive plate chambers (RPCs), the thin gap chambers (TGCs) and the NSW can trigger single muon signatures.

4.2.4 TRIGGER AND DATA ACQUISITION

The trigger and data acquisition (TDAQ) system selects events that are interesting for physics analyses according to the physics goals of ATLAS. The system, schematically shown in fig. 4.9, is split into three parts: the hardware-based level-1 (L1) trigger system, the software-based high-level trigger (HLT) system, and the actual acquisition (DAQ) system. The custom electronics of the L1 trigger combine reduced granularity information from the muon and calorimeter detectors to decide within $2.5\mu\text{s}$ whether an event is interesting. Due to rate restrictions by the detector read-out systems, the collision rate of 40 MHz must be reduced to 100 kHz. Events with positive L1 trigger decisions are then passed on to the read-out system. The trigger rate is shared between the different physics objects considered attractive by the ATLAS collaboration at the time. The following HLT trigger stage is purely software-based and runs on commercial computers. The HLT rate is limited to 3 kHz by data storage possibilities for offline reconstruction; a typical data rate is 6 GB/s. A precise reconstruction of the physics objects would be computationally too expensive. However, the HLT software is designed to mimic the full offline physics object reconstruction as closely as possible. After a favourable HLT decision, the data is stored for offline reconstruction and analysis. In case of CERN mass storage services downtime, a storage volume of 1.4 PB is available close to the ATLAS detector to guarantee the storage of 24 h of data taking.




Figure 4.9: Overview of the trigger and data acquisition system of the ATLAS detector. The level-1 (L1) and high-level trigger (HLT) systems deliver their decisions to the detector read-out and data flow scheme on the right that transmits and stores the collision data. The 40 MHz collision rate is ultimately reduced to a 3 kHz output rate of 6 GB/s that stores events interesting for physics analyses. Image taken from ref. [104].

4.3 PHYSICS OBJECT RECONSTRUCTION

Involved algorithms embedded in the ATLAS software framework *ATHENA* [106], named after the ancient Greek goddess of wisdom, make sense of the multitude of hits from the different subdetectors. Figure 4.10 illustrates where different types of particles deposit their energy and which part of the detector is responsible for their detection. The inner detector (pixel, SCT, and TRT) reconstructs the charged particle trajectories. Photons and electrons leave all their energy in the electromagnetic calorimeter, while hadrons, such as pions, protons, and neutrons, make it into the hadronic calorimeter before leaving a hadronic shower signature as illustrated in fig. 4.10. Muons can travel through the complete detector, leaving minimal ionising energy signatures in the calorimeter and the muon spectrometer. Neutrinos are invisible to the ATLAS detector. The following sections describe the reconstruction algorithms required to rebuild the physics objects needed for this work: tracks (section 4.3.1), primary vertices (section 4.3.2), muons (section 4.3.3), electrons (section 4.3.4), large-radius tracks (section 4.4.1), and displaced vertices (section 4.4.2).

4.3.1 TRACKS

Tracks are defined in ATLAS via their impact parameters (d_0 , z_0), their azimuthal and polar angles (ϕ , Θ), and their charge over momentum ratio (q/p). Figure 4.11 defines those parameters in the ATLAS coordinate




Figure 4.10: Schematic overview of particle candidates and where they are reconstructed in the ATLAS detector. The scheme shows a cut in the $x - y$ plane of the detector. Image taken from ref. [107].




Figure 4.11: Schematic illustration defining the track parameters used in ATLAS. A cylindric coordinate system is used with the origin at the nominal interaction point at the centre of the detector, with the z -axis pointing along the beam direction, the x -axis towards the centre of the LHC ring, and the y -axis upward. Shown are the two impact parameters d_0 and z_0 , the azimuthal angle ϕ around the beam axis, the polar angle Θ that is also expressed as pseudorapidity $\eta = -\ln \tan(\Theta/2)$, the transverse momentum p_T defined in the $x - y$ plane, and the momentum vector $\mathbf{p}$. Image taken from ref. [108].




Figure 4.12: Illustration of (a) single-particle clusters on a pixel sensor and (b) a merged pixel cluster due to very collimated charged particles. Different colours represent energy deposits from different charged particles traversing the sensor, and the particle trajectories are shown as arrows. Figure and caption are taken from ref. [109].

system. The impact parameter d_0 marks the closest distance of the track to the beam line or z axis, while z_0 marks the corresponding shift on the z axis with respect to the origin of the coordinate system.

In the ATLAS track reconstruction, tracks are formed first in the silicon part of the inner detector [109] (pixel and SCT) and then extended to the TRT in an inside-out and outside-in tracking scheme [110]. The primary track reconstruction in the pixel and SCT detectors starts with the clusterisation part of the raw measurements in the detector modules. A so-called connected compound analysis algorithm [111] groups charge measurements above a given threshold into clusters and connects cluster groups from the pixel and SCT subdetectors. Space points, also called hits, are three-dimensional measurements that are formed from clusters where the charged particles cross the active detector area. In the pixel detector, each cluster is one space point, while in the SCT detector, clusters from both sides of the strip layer need to be combined to receive a three-dimensional measurement. The charge of each cluster is determined with a linear interpolation technique further described in ref. [112]. Using so-called Monte Carlo (MC) simulation 'truth information'² at the MC event generator level and reconstructed quantities, the clusters can be classified as whether they accompany only one (single particle cluster) or multiple particle trajectories (merged cluster). Figure 4.12 portrays the cluster classification schematically. In fig. 4.12 (a), the spatial separation of the tracks indicated as arrows is high enough compared to the pixel density to form a single-particle pixel cluster per track. In contrast, fig. 4.12 (b) shows a merged pixel cluster that contains three very collimated tracks, indicated as arrows. Multiple tracks share so-called shared clusters. Those shared clusters indicate incorrect assignments, and their maximum number per track is limited in the track reconstruction process.

²The so-called Monte Carlo simulation 'truth information' contains information about the origin of all generated particles and their links. If, for example, particles decay into other particles, those connections are stored. Comparing this information with the reconstructed quantities, such as energy clusters, helps us to understand the capabilities of the reconstruction algorithms.




Figure 4.13: A sample event from a $t\bar{t}$ MC simulation showing the two possible TRT hit associations (only the barrel measurements): the brighter coloured hits show the extensions that follow the track found by seeds in the silicon, while the darker grey ones show TRT hits found through the stand-alone track segment search. The TRT segments are then used for backtracking into the silicon to optimise the overall reconstruction efficiency. Figure and modified caption are taken from ref. [110].

In the next step, track seeds are formed from three space points, whose impact parameter is estimated assuming a helical trajectory in a uniform magnetic field. The best quality seeds are mixed detector seeds that include space points from both pixel and SCT detectors, followed by pixel-only and SCT-only seeds. Track candidates are created using a combinatorial Kalman filter [113] that adds additional space points. Multiple track candidates per seed are allowed in applying the Kalman filter. To be considered a track candidate, the track seed must be at least compatible with an additional space point. A track score system is used to rank tracks in terms of their quality. Each assigned cluster increases the track score, while holes decrease the track score. Holes are detector elements that are crossed by the trajectory of the track candidate without forming a cluster. In addition, the χ^2 goodness of fit value from the track fit is included in the track score, and high p_T tracks are favoured over low p_T tracks.

With the track scores in hand, the ambiguity solver can disentangle the situation of clusters assigned to multiple track candidates (shared clusters). The guiding principles are that no more than two tracks can share a cluster, and each track can share a maximum of two shared clusters. If tracks fail the shared cluster criteria, the algorithm can remove clusters from the track and assign a new score to see if the modified track would make it to the next step.

In summary, track candidates must fulfil the following criteria: $p_T > 400$ MeV, $|\eta| < 2.5$, ≥ 7 pixel and SCT hits, maximal one shared pixel cluster or two shared SCT clusters on the same layer, ≤ 2 holes for pixel and SCT, ≤ 1 hole for pixel, $|d_0| < 2$ mm, and $|z_0 \sin \Theta| < 3$ mm. In the pixel detector, a neural network is used when track candidates are assigned to multiple clusters to identify merged clusters in the pixel detector using the charge and relative position of the pixels [114]. The following final track fit uses all available information, and track candidates that pass the ambiguity solver step are added to the final track collection. Each cluster has an assigned position and uncertainty determined by a second neural network [114]. Tracking efficiency plots for Run2 are given in ref. [109].

Now that we have track candidates from the silicon part of the inner detector, we can move on and try to extend them to the TRT. Two schemes are used here: the actual extension algorithm and a processing and evaluation part [110]. A deterministic annealing filter [115], which is an extension to the standard Kalman filter that is optimised for high hit densities, is used to extend the track candidate to the TRT. Figure 4.13 shows the matched TRT extensions together with the silicon tracks in red. In a second approach, the tracking is reversed from the outside to the inside of the inner detector. Leftover hits in the TRT are grouped into segments by a global pattern recognition algorithm [110] and traced back into the silicon tracker to see if silicon segments have been missed previously. The TRT segments are also shown in fig. 4.13 as dark grey dots. After successful TRT extension, the tracks are refitted and added to the final INDETRACKPARTICLES track collection that can be used on the analysis level.




Figure 4.14: Illustrative event display without dimensions of multiple interactions along the beamline (pile-up). The reconstructed primary vertices are visible as coloured squares. The associated tracks are shown as lines in the same colour. The mean proton bunch length at the interaction point is 7.94 cm, assuming a Gaussian beam distribution [91]. Image taken from ref. [116] (Run number: 201289, event number: 24151616).

4.3.2 PRIMARY VERTICES, PILEUP, AND LUMINOSITY

At the ATLAS interaction point, multiple pp interactions occur along the beamline in the same bunch crossing (as shown in fig. 4.14). A precise primary vertex reconstruction is required to disentangle the multitude of tracks from the various interactions. Besides the interesting hard-scatter vertex³, we see many additional vertices that usually originate from additional pp interactions with low transverse momentum transfer. This superposition of vertices is referred to as pile-up. The amount of pile-up is quantified by the average number of inelastic pp interactions per bunch crossing, μ , under constant beam conditions. μ changes over time, depending on the beam intensity, and is directly related to the instantaneous luminosity $\mathcal{L}$ [117, 118]. The pp interaction rate dN/dt is equal to the luminosity times the production cross-section σ_{pp} [119]:

$$\frac{dN}{dt} = \mathcal{L} \cdot \sigma_{pp} \quad (4.3)$$

The instantaneous luminosity for a single colliding bunch pair $\mathcal{L}_b$ is proportional to the revolution frequency f_r of the LHC, which is 11246 Hz for protons times the average number of pp interactions per bunch crossing μ divided by the cross-section that corresponds to inelastic pp collisions σ_{inel} .

$$\mathcal{L}_b = \frac{\mu f_r}{\sigma_{\text{inel}}} = \frac{\mu_{\text{vis}} f_r}{\sigma_{\text{vis}}} \quad (4.4)$$

μ_{vis} is the visible interaction rate per bunch crossing, and σ_{vis} corresponds to the part of σ_{inel} that can be reconstructed in the detector. σ_{vis} is a calibration constant that can be determined using so-called Van der Meer scans (for details, see ref. [118]). The quantity used at the analysis level is the integrated luminosity L ,

³The so-called hard-scatter vertex describes the part of the pp interaction that contributes to the main event of interest, with a high transverse momentum transfer in the scattering process.




Figure 4.15: a) Luminosity weighted distribution of the mean number of interactions per bunch crossing μ for the 2015-2018 pp collision data. b) Cumulative luminosity as a function of time that is delivered to ATLAS (green), recorded by ATLAS (yellow), and certified to be good for physics analysis (blue). Figures taken from ref. [120].

which follows from integrating $\mathcal{L}$ over time for a particular time interval, for example, 2015-2018, which would be the integrated luminosity that corresponds to the full LHC Run2.

$$L = \int \mathcal{L}(t) dt \quad \mathcal{L}(t) = \sum_{b \in \text{bunches}} \mathcal{L}_b(t) \quad (4.5)$$

The total instantaneous luminosity $\mathcal{L}$ is obtained by summing over all colliding bunches b . Figure 4.15 shows the mean number of interactions per bunch crossing μ for the different years of data taking during LHC Run2 and the cumulative distribution of the integrated luminosity L for the same period. ATLAS has been unable to record approximately 9 fb^{-1} of data due to downtimes in data acquisition. A total integrated luminosity of $L = 139 \text{ fb}^{-1}$ is considered suitable for physics analyses.

The reconstruction of the primary vertex starts with the reconstructed track collection and follows two steps: the vertex finding and the vertex fitting [121]. All vertices with at least two tracks are primary vertex candidates. The output of the vertex fitter is a set of three-dimensional vertex positions and their corresponding covariance matrices. In each event, the vertex with the highest $\sum_{\text{tracks}} p_T^2$ is considered as the primary vertex [122].

4.3.3 MUONS

The reconstruction of muons splits into two categories: standalone muons that only use MS information in their reconstruction and muons that are reconstructed using all available information from the subdetector systems [123]. Typical muon origins at the LHC are W and Z decays and decays from heavy flavour hadrons. Those decays happen very close to the pp interaction point.

For standalone muons, the reconstruction starts with forming so-called segments, which are straight-line approximations of track segments in an individual MS station. The segments are identified from the individual hits using a Hough transformation [124]. Segments of the different MS stations are then combined into track candidates that point loosely to the interaction point. A simple parabolic trajectory is used as an approximation of the bending in the magnetic field. The final global χ^2 fit for the muon trajectory includes effects of material interactions and misalignments between the detector chambers. Outlier hits that don't fit the candidate trajectory are removed in an additional track optimisation step, while other potentially matching hits are added. The track fit is then repeated, including the new set of hits. Tracks that share many hits with higher quality tracks are removed to solve ambiguities. Finally, the track candidates are extrapolated to the beamline and refitted while considering the muon's energy loss in the calorimeters.

TYPES OF MUONS

Muon candidates that are reconstructed using all subdetector systems are classified into several groups: combined (CB), inside-out combined (IO), muon spectrometer extrapolated (ME), segment tagged (ST), and calorimeter tagged (CT).

COMBINED MUONS The guiding principle for combined muons is to match the MS tracks to the ID tracks. A combined fit is performed that is based on all hits in the ID and the MS, and that takes the energy losses in the calorimeters into account. Similar to the standalone muon reconstruction, a track optimisation step can update the set of hits and the track fit is repeated.

INSIDE-OUT-COMBINED The inside-out muon reconstruction scheme tries to extrapolate ID track candidates into the MS and looks for at least three loosely aligned MS hits. The combined fit includes all hits and the energy losses in the calorimeters. IO muons recover some reconstruction efficiency because the technique does not depend on independently reconstructed MS tracks and is interesting for low p_T muons that potentially only reach the first layer of MS.

MUON SPECTROMETER EXTRAPOLATED The MS track parameters are extrapolated to the beamline in cases where the ID track matching with the MS tracks fails. The additional ME reconstruction branch extends the acceptance for muons that are created outside the ID, such as muons from pion or kaon decays.

SEGMENT TAGGED Here, the ID tracks are extrapolated to the MS as well, but they fail the requirement of IO muons. The ST muons are matched with at least one MS segment. The ID track parameters are taken as the ST muon parameters.

CALORIMETER TAGGED Muons leave characteristic energy deposits in the calorimeters. Those characteristics are utilised to reconstruct CT muons, where the ID track is only extrapolated to the calorimeters.

The muon parameters are again taken from the ID track parameters. Here, larger contributions from tracks misidentified as muons are expected.

WORKING POINT DEFINITION

Now that a set of different muon candidates is formed, we can order them regarding their "muonness". The muon quality depends on various parameters, such as the number of hits in the ID subdetectors, the number of MS stations involved in the MS track formation, the track fit properties themselves, and variables that quantify the compatibility between the subsystems. The particular set of parameters applied to the muon candidate is called the muon working point. Three standard working points are defined in ATLAS, listed with increasing purity and decreasing identification efficiency: `LOOSE`, `MEDIUM`, and `TIGHT`. The working points are inclusive in a way that all `TIGHT` muons also fulfil the `MEDIUM` working point, and all `MEDIUM` muons also fulfil the `LOOSE` working point. Muon reconstruction efficiencies for the three working points as a function of transverse momentum and pseudorapidity are shown in fig. 4.16.

To define the working points, we need to start with their defining parameters first. Precision stations are MS stations in which the muon has at least three hits in the MDT or CSC detectors. Precision hole stations are MS stations where muons have less than three hits and miss at least three hits that are expected by the trajectory. The charge over momentum q/p compatibility is defined only for CB and IO muons. It depends on the track parameters from the ID and the MS $q/p_{\text{ID,MS}}$ and their uncertainties $\sigma(q/p_{\text{ID,MS}})$:

$$q/p = \frac{|q/p_{\text{ID}} - q/p_{\text{MS}}|}{\sqrt{\sigma(q/p_{\text{ID}})^2 + \sigma(q/p_{\text{MS}})^2}} \quad (4.6)$$

The `MEDIUM` working point for muons, which is used exclusively in this thesis, then requires ≥ 1 pixel hit, ≥ 5 SCT hits, ≤ 2 holes in pixel and SCT for CB, IO, ST, and CT muons. CB and IO muons need ≥ 2 precision MS stations (≥ 1 precision MS station if ≥ 1 precision hole MS station) for $|\eta| < 2.5$ ($|\eta| < 0.1$) and $q/p < 7$. Definitions of the `LOOSE` and `TIGHT` working points can be found in ref. [123].

4.3.4 ELECTRONS

The electron reconstruction in ATLAS can be split into three steps: the seed cluster reconstruction, the track reconstruction, and the actual electron candidate reconstruction [125]. When high-energy electrons come into contact with the detector material, they lose energy primarily due to bremsstrahlung. The photons produced can then convert into electron-positron pairs that interact with the detector material again. Those electrons, positrons and photons are emitted in a collimated way and typically reconstructed as part of the same electromagnetic cluster. This shower can also happen in the inner detector or at the beam pipe, resulting in numerous tracks. For a successful electron reconstruction, a few key ingredients are needed: localised clusters of energy deposits in the electromagnetic (EM) calorimeter, charged particle tracks in the inner detector, and




Figure 4.16: Muon reconstruction and identification efficiencies for the LOOSE, MEDIUM, and TIGHT working points. The left plot shows the efficiencies measured in $J/\psi \rightarrow \mu^+\mu^-$ events as a function of p_T . The right plot displays the efficiencies measured in $Z \rightarrow \mu^+\mu^-$ events as a function of η , for muons with $p_T > 10$ GeV. The predicted efficiencies are depicted as open markers, while filled markers illustrate the result of the measurement in collision data. The statistical uncertainty in the efficiency measurement is smaller than the size of the markers, and thus not displayed. The panel at the bottom shows the ratio of the measured to predicted efficiencies, with statistical and systematic uncertainties. Figures are taken from ref. [123].

a close matching in the $\eta \times \phi$ space of the tracks to the clusters formed in the EM calorimeter. Figure 4.17 schematically shows the parts of the ATLAS detector involved in the reconstruction.

The energy cluster candidates in the EM calorimeter are seeded from localised energy deposits using a sliding window algorithm if the transverse energy exceeds 2.5 GeV. The reconstruction efficiency of the seed cluster formation ranges from 65% at $E_T = 4.5$ GeV to 99% at $E_T = 15$ GeV [125].

The track reconstruction [109] for electron candidates is done in two ways: the standard ATLAS track reconstruction following the ATLAS Global χ^2 Track Fitter [126] and a fitting procedure using a Gaussian-sum filter (GSF) [127] that is optimised for electron tracks and that shows better performance in dealing with energy losses of charged particles in the detector material. The GSF method [128] is a generalisation of the Kalman filter scheme [113] that is usually used in tracking. Non-linear effects related to bremsstrahlung are assessed by the GSF method, and a sum of Gaussian functions models experimental noise. Multiple Kalman filters run in parallel, and a weighted sum of Gaussian functions approximates each track parameter.

Finally, the GSF track candidate is matched to the seed cluster in the EM calorimeter. Electron candidates must have no association with a vertex from photon conversion [129]. The final calibration of the EM clusters to obtain the original electron energy is performed using multivariate techniques [130].

LIKELIHOOD ID SCHEME Instead of cutting on several variables, a likelihood identification scheme [125] is used to identify electron candidates. A likelihood-based identification is more inclusive than a cut-based one: candidates that would fail the selection because of one variable could still pass the likelihood-based identification.




Figure 4.17: 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. Figure and modified caption are taken from [125].




Figure 4.18: The transformed LH-based identification discriminant d'_L for reconstructed electron candidates with good quality tracks with $30 \text{ GeV} < E_T < 35 \text{ GeV}$ and $|\eta| < 0.6$. The black histogram is for prompt electrons in a $Z \rightarrow e^+e^-$ simulation sample, and the red (dashed-line) histogram is for backgrounds in a generic two-to-two process simulation sample. The histograms are normalised to unit area. Figure and caption are taken from [125].




Figure 4.19: Identification efficiencies of electrons from $Z \rightarrow e^+e^-$ decays as a function of the electrons (a) transverse momentum and (b) pseudorapidity for the different identification working points. The top panels show the efficiencies obtained in data and simulation, with their statistical and total uncertainties displayed as inner and outer error bars. The middle panels show their ratio applied as correction factors in analyses. The bottom panels show the statistical and the total uncertainties in the data/MC ratio. Figures are taken from ref. [131].

Two likelihoods are defined for signal S and background B :

$$L_S(\vec{x}) = \prod_{i=1}^n P_{S_i}(x_i) \quad \text{and} \quad L_B(\vec{x}) = \prod_{i=1}^n P_{B_i}(x_i) \quad (4.7)$$

Each likelihood comprises a product of n probability density functions P_{S_i/B_i} for each input parameter x_i . The likelihoods are formed by input parameters such as the energy ratio of the first layer in the hadronic calorimeter and the electromagnetic calorimeter (also called hadronic leakage), energy ratios within layers of the electromagnetic calorimeter, track properties such as number of hits, transverse impact parameter d_0 , significance of the transverse impact parameter (defined as the ratio of d_0 to its uncertainty: $|d_0/\sigma(d_0)|$), and momentum difference between the closest point to the beam axis and last measurement point. Details about the input parameter x_i are given in [125].

For each electron candidate, a discriminant is formed:

$$d_L = \frac{L_S}{L_S + L_B} \quad (4.8)$$

d_L peaks at 1 for the signal and at 0 for the background. This discriminant is further transformed to simplify a working point (WP) definition:

$$d'_L = -\frac{1}{15} \ln\left(\frac{1}{d_L} - 1\right) \quad (4.9)$$

WPs are defined based on numeric cutoffs on d'_L , such that electron candidates that exceed a particular value of d'_L are considered as signals and, therefore, electrons. Figure 4.18 shows the discrimination power of the transformed discriminant d'_L between prompt electrons from a $Z \rightarrow ee$ simulation sample and background from multijet production, W and Z production, and events featuring top quarks.

Standard inclusive WPs are defined that correspond to different values of d'_L . The WPs are LHVERY-LOOSE, LHLOOSE, LHMEDIUM, and LHTIGHT. Inclusive means that electron candidates that pass the LHTIGHT WP also pass all looser WPs. Depending on the WP, there are different requirements for the number of hits in the inner detector:

- LHTIGHT, LHMEDIUM, LHLOOSE: ≥ 2 pixel hits + ≥ 7 combined pixel and SCT hits.
- LHTIGHT, LHMEDIUM: one of the pixel hits must be in the innermost functional pixel layer (helps to reduce background from photon conversions).
- LHVERYLOOSE: ≥ 1 hit in the pixel detector.

Figure 4.19 summarises the electron identification efficiencies for the three different working points LHTIGHT, LHMEDIUM, and LHLOOSE as a function of transverse energy E_T and pseudorapidity η . As expected, the loosest WP reaches the highest identification efficiency, as the selection is less strict. The efficiency of electron identification is lower for lower transverse energies and reaches its maximum for $E_T > 80$ GeV. For further details, see ref. [131].

Specific to this work, we define an additional WP to increase the reconstruction efficiency for highly displaced electron candidates with low transverse energy relevant for long-lived particle decays. An ATLAS search has already used this additional WP for displaced leptons; see ref. [132] for details on the improvement of the reconstruction efficiency. Ref. [132] quotes a reconstruction and identification efficiency of 80% for electrons with $|d_0| < 10$ mm and $p_T > 20$ GeV. We call the WP LHVERYVERYLOOSE since it removes all requirements on the impact parameter d_0 and imposes no minimum number of hits for the pixel or SCT detector. Other than that, it is identical to LHVERYLOOSE. However, there are requirements on displaced vertex tracks that will be discussed later in this chapter.

Table 4.1: Listing of the main differences between standard and large-radius tracking [133].

	standard	large-radius tracking
$\max(d_0)/\text{mm}$	10	300
$\max(z_0)/\text{mm}$	250	1500
$\max(\eta)$	2.7	5
$\max(\# \text{ of shared silicon clusters})$	1	2
$\min(\# \text{ of unshared silicon hits})$	6	5
Kalman filter type	combinatorial	sequential

4.4 SPECIAL LONG-LIVED PARTICLE RECONSTRUCTION

A dedicated reconstruction chain is necessary to reconstruct signatures of long-lived particles in the ATLAS detector. The two main reconstruction methods used are large-radius tracking and secondary vertexing. Both are briefly outlined below.

4.4.1 LARGE-RADIUS TRACKING

The standard tracking schemes pursued in the ATLAS reconstruction chain assume that particles are produced close to the interaction point (IP). This is why dedicated long-lived particle searches must rely on a third type of tracking next to the standard inside-out and outside-in algorithms described in section 4.3.1: the so-called large-radius tracking (LRT) [133]. LRT substantially widens the track parameter and loosens some requirements compared to standard tracking to incorporate that decay products of long-lived particles emerge further away from the IP and retain a high reconstruction efficiency level. Table 4.1 highlights the differences between the standard tracking and LRT.

LRT follows the strategy of standard inside-out tracking closely. Track seeds are created from leftover hits that the standard tracking algorithms have not used. Then, a pattern recognition algorithm creates track seeds followed by a window search that forms the track candidate. The sequential Kalman filter [113] used here is one key difference from the standard tracking since it forms only one possible track candidate from the hit pattern. The following ambiguity solver step uses the relaxed parameters listed in table 4.1. Track candidates that make it through the ambiguity solver step are, if possible, extended to the TRT exactly as it is done in the standard tracking. Tracks with a TRT extension are added to the final LRT track collection and merged into the standard track collection. Figure 4.20 shows the reconstruction efficiency for displaced leptonic and hadronic tracks as a function of the radius r_{prod} , with r_{prod} being the distance in the $x - y$ plane between the beamline and the point where the displaced tracks emerge. For both leptonic and hadronic tracks, the reconstruction efficiency decreases with increasing r_{prod} . For high r_{prod} values, LRT can reconstruct displaced tracks where the standard tracking fails. Unfortunately, even the combined reconstruction efficiency of standard and large-radius tracks decreases for increasing r_{prod} . Reasons for this are algorithmic inefficiencies and the fact that the ATLAS detector is optimised for prompt tracks, not largely displaced ones [133].




Figure 4.20: Inclusive track reconstruction efficiency for displaced charged particles produced by the decay of long-lived signal particles. The efficiency is shown as a function of the radius of the production of displaced particles r_{prod} (defined in the $x - y$ plane between the beamline and the production point of the displaced particles), with truth particles and tracks subject to fiducial selections. The efficiency for standard and large-radius tracking is additive, and the total efficiency is the sum of both. Figures and modified are taken from ref. [133].

4.4.2 DISPLACED VERTEX RECONSTRUCTION

Now that we have reconstructed large-radius tracks, we can continue and describe how the VertexSecondary-Inclusive (VSI) algorithm [134] is used to reconstruct displaced vertices from long-lived particle decays with macroscopic lifetimes $c\tau \geq \mathcal{O}(1 \text{ mm})$. The algorithm is designed to be as flexible as possible to allow fine-tuning for a wide range of long-lived particle signatures.

In the first step, two-track seed vertices are formed. However, not all tracks can serve as input tracks due to the high track multiplicity. Seeding tracks need to have: $p_T > 1 \text{ GeV}$, no association to a primary vertex from hard-scatter or a vertex from pileup interactions [122], ≥ 6 SCT hits for tracks with zero pixel hits, ≥ 1 TRT hit for tracks with ≤ 2 pixel hits, and ≥ 7 SCT hits for tracks with $p_T < 20 \text{ GeV}$, ≥ 20 TRT hits for tracks with $p_T < 20 \text{ GeV}$ and $|\eta| < 1.5$, and $|d_0| > 2 \text{ mm}$. Two-track seed vertices are formed from all possible combinations of these preselected tracks. A χ^2 vertex fit is performed [135] and the resulting vertex position r_{vtx} needs to be smaller than the SCT/TRT boundary at 563 mm. Both tracks that form the vertex candidate need to pass a hit pattern consistency check that consists of allowed, required, and forbidden hits as illustrated in fig. 4.21.

At this stage, a single track can be used multiple times in different seed vertices. An undirected graph network is used to form multi-track vertices [134]. Now, only the association with the best-fitting vertex (smallest $\chi^2_{\text{vtx}}/n_{\text{dof}}$, with the degrees of freedom $n_{\text{dof}} = 2n - 3$ for a n -track vertex [134]) is kept for each track.

In the vertex merging step, nearby vertices are merged into a single multitrack vertex. The vertex fit is repeated after the merge.

The last step in the vertex reconstruction is the track attachment step. Here, nearby tracks can be attached to vertex candidates. The tracks considered in the track attachment step are only required to pass $p_T > 1 \text{ GeV}$,




Figure 4.21: The schematic shows the transverse view of the pixel detector in the barrel region. (a) The vertex is between the two layers of B-Layer and Layer-1 sensors. The tracks of the reconstructed secondary vertex must not have hits on the layers within the vertex radius (i.e. IBL and B-Layer) and must have hits on the closest layer outside the vertex (i.e. Layer-1). (b) The vertex is close to the pixel B-Layer sensors. In this example, where the vertex is inside the B-Layer, the tracks are not allowed to have hits on the IBL (forbidden hits, blue crosses) but may have hits on the B-Layer (allowed hits, black circles) and must have hits on the Layer-1 (required hits, red dots). Analogous requirements are made on vertices close to the other layers. Figures and modified caption are taken from ref. [134].

and tracks that are already associated with primary pp interactions are included. The vertex fit is repeated for each track attachment.

Now, a final fit is performed, including all successfully attached tracks. The invariant mass of the vertex is calculated assuming $m_{\pi^\pm}$ for all tracks. Figure 4.22 shows the vertex reconstruction efficiency and the difference between the true and reconstructed vertex position for different types of long-lived particle decays (for details, see ref. [134]). The reconstruction efficiency is lower for vertices with lower track multiplicity. The vertex position resolution depends on the resolution of the track parameters, but fig. 4.22 shows good agreement between the true $r_{\text{LLP decay}}^{\text{truth}}$ and reconstructed vertex positions r_{reco} .

As it will be discussed in section 5.5.1, we optimised the VSI algorithm for the long-lived signature of leptonic HNL decays with their low track multiplicity to reach higher vertexing efficiencies while keeping the background at a manageable level.





Figure 4.22: Shown here is the (a) reconstruction efficiency and (b) the difference between the true (MC event generator level) and reconstructed decay position for different long-lived particle decays. Figures taken from ref. [134].

4.5 SIMULATION

It is critical to have precise simulations of the signature that HNLs leave in the detector to understand their kinematics and to design the criteria to select events that could contain HNLs. Also, possible SM background processes need to be simulated to predict the number of expected events that leave similar traces in the detector. The matrix elements $\mathcal{M}$ of the processes of interest are too complex to compute beyond the first few orders of perturbation theory. Due to the complexity of the processes and the necessity to integrate over a final state phase space of huge and variable dimensions, those simulations rely on numerical Monte Carlo (MC) integration methods [138]. This section briefly summarises the simulation methods used at the LHC, particularly in the ATLAS experiment.

In the simulation, the various subprocesses are factorised into different momentum scales to simplify their treatment. Figure 4.23 illustrates the subprocesses involved in the pp interaction. The red blob in the centre represents the hard collision (or 'hard-scatter') part of the simulation, where the constituent partons of the incoming beams interact and produce a small number of highly energetic outgoing partons, leptons, or gauge bosons. The actual event of interest is simulated in the hard-scatter part of the simulation. Due to the asymptotic freedom of QCD, the high-energy part of the simulation that involves a high momentum transfer can be solved using perturbation theory, with the corresponding matrix element being computed e.g. via Feynman diagrams.

The cross-section $\sigma_{ab \rightarrow n}$ of the hard subprocess $ab \rightarrow n$ with the initial state partons a, b and the final state particles n can be computed using a factorisation formula [139, 140]:

$$\sigma_{ab \rightarrow n} = \sum_{a,b} \int_0^1 dx_a dx_b \int d\phi_n f_a^{h_1}(x_a, \mu_F) f_b^{h_2}(x_b, \mu_F) \frac{1}{2x_a x_b s} |\mathcal{M}_{ab \rightarrow n}|^2(\phi_n; \mu_F, \mu_R) \quad (4.10)$$

where $f_{a/b}^{h_1/h_2}(x_{a/b}, \mu_F)$ are the parton distribution functions (PDFs) that depend on the momentum fraction $x_{a/b}$ of parton a/b to its parent hadron h_1/h_2 momentum and the so-called factorisation scale μ_F . The matrix element squared $|\mathcal{M}_{ab \rightarrow n}|^2(\phi_n; \mu_F, \mu_R)$ follows from the sum of possible Feynman diagrams up to a particular order in perturbation theory. The choice of the PDF set is usually connected to the choice of the factorisation and renormalisation scale μ_R . A single energy scale Q^2 can be defined such that $\mu_F = \mu_R = Q^2$. Examples of proton PDFs for $Q^2 = 10 \text{ GeV}^2$ and $Q^2 = 10^4 \text{ GeV}^2$ are shown in fig. 4.24.

s in eq. (4.10) is the centre-of-mass energy squared of the two colliding protons and $d\phi_n$ is the differential phase space element for the n final state particles:

$$d\phi_n = \prod_{i=1}^n \frac{d^3 p_i}{(2\pi)^3 2E_i} \cdot (2\pi)^4 \delta^{(4)} \left(x_a P_a + x_b P_b - \sum_{i=1}^n p_i \right) \quad (4.11)$$

with the momenta p_i and energies E_i of the final state particles and the initial state hadron momenta $P_{a/b}$. $\delta^{(4)}$ is the Kronecker-Delta function in four dimensions.




Figure 4.23: Sketch of a hadron-hadron collision as simulated by a Monte Carlo event generator. The red blob in the centre represents the hard collision, surrounded by a tree-like structure representing Bremsstrahlung as simulated by parton showers. The purple blob indicates a secondary hard scattering event. Parton-to-hadron transitions are represented by light green blobs, dark green blobs indicate hadron decay, and yellow lines signal soft (low p_T) photon radiation. Figure and modified caption are taken from ref. [136].




Figure 4.24: Example of MMHT2014 NNLO parton distribution functions (PDFs) for protons (with valence quark content of two up quarks and one down quark) at an energy scale $Q^2 = 10 \text{ GeV}^2$ (left) and $Q^2 = 10^4 \text{ GeV}^2$ (right) with associated 68 % confidence-level uncertainty bands. $f(x, Q^2)$ is the probability of finding a parton with a momentum fraction x from the parent proton when the proton is probed at the energy scale Q^2 . Shown are PDFs for up and down valence quarks (u_V, d_V) that dominate at high x and PDFs for gluons (g), anti-up and anti-down quarks ($\bar{u}, \bar{d}$) as well as strange and charm quarks (s, c) that dominate at low x . The probability $f(x, Q^2)$ is scaled by x on the vertical axis. Figure and modified caption are taken from ref. [137].

Now, continuing outwards from the hard-scatter centre in fig. 4.23, we see various other subprocesses involved as we go down in energy. The tree-like structure in fig. 4.23 represents the parton showers of incoming and outgoing coloured participants in the subprocess. The partons can radiate gluons, while the gluons can radiate other gluons or produce pairs of quarks and antiquarks, which causes a shower of outgoing partons. This parton shower is likely collinear with the outgoing partons and forms a structure of a localised bundle that contains most of the energy of the outgoing parton, also called a jet. Another subprocess is the secondary interaction, often called the underlying event, which describes the interaction of the remaining partons of the colliding hadrons that are not directly involved in the hard-scatter process (indicated by a purple blob in fig. 4.23). Non-perturbative strong interactions that convert the partons into outgoing hadrons and connect them to the incoming beam hadrons indicated as light green blobs in fig. 4.23 happen at low energy scales ($\mathcal{O}(1) \text{ GeV}$). In this process called hadronisation, phenomenological models inspired by QCD are used, as the strong coupling constant in the low energy regime is too large for the use of perturbation theory. The two main models here are the string model [141] and the cluster model [142]. The set of values for the parameters (for a given model) obtained from a fit to data used for parton showering, hadronisation and the underlying event is called 'tune' and corresponds to the used PDF set. The models are predictive for new collision types or energies once tuned to a dataset. The last important subprocess illustrated as dark green blobs in fig. 4.23 are decays of hadrons with a lifetime smaller than the detector's size.

In ATLAS, everything related to the MC simulation is embedded in the ATLAS simulation infrastructure [143]. The full event simulation has to be comparable with real detector data from the pp collisions, such that both real and simulated data end up in the same data format that the reconstruction algorithms can process. Once the event is simulated using an event generator such as PYTHIA 8.212 [144], which can simulate the hard-scatter event itself and the parton showering, we still need to take care of the detector response that includes the reconstruction of the physics objects, as discussed in this chapter. The particles are propagated through a detailed three-dimensional model of the ATLAS detector using the GEANT4 simulation [145]. Afterwards, the energy deposits created by the particles in the different subdetectors are digitised. This digitised energy deposit can be binary in a submodule of the detector, where a voltage or current rises above a predefined threshold, creating a hit. Moreover, some subdetectors store the digitised version of the full signal current or voltage detector waveform as a function of time. Every subdetector is equipped with a unique digitisation software. After the digitisation step, the simulated and real detector data can be run through the same trigger and reconstruction frameworks. The simulation framework also incorporates effects related to the detector and pile-up conditions at a particular time of data-taking. The simulated events are scaled in a process called pile-up reweighting to reproduce the distribution of the average number of interactions per bunch crossing μ observed in the data. Depending on the use case and the number of events required to be simulated, the CPU-intense full simulation of the ATLAS detector response can be accelerated by one of the so-called fast-simulation schemes available in ATLAS. For example, the ATLAS fast (Atlfast-II) simulation [146] saves about 80 % of the simulation time spent with hadron showers in GEANT4 using pre-simulated showers that can be restored from libraries. The ATLAS simulation infrastructure also entails the validation of the simulation output against known physics processes, as well as software validation and performance testing. The simulated samples used in this work will be discussed in section 5.4.1.

PART II

DATA ANALYSIS

5 SEARCH FOR LONG-LIVED HEAVY NEUTRAL LEPTONS

In 2023, the search results for long-lived HNLs using the ATLAS detector and the full LHC Run2 dataset (2015-2018), including 139 fb^{-1} of 13 TeV pp collision data, were published [10]. This chapter describes the data analysis and its key findings.

5.1 SENSITIVITY

HNLs ($\mathcal{N}$) are typically parameterised by their squared couplings to SM neutrinos $|U|^2$ and their mass $m_{\mathcal{N}}$. Experiments then simplify the HNL parameterisation by assuming only one HNL that couples to only one SM neutrino flavour $|U_{\alpha}|^2$, with $\alpha = e, \mu, \tau$. Different experiments can probe different regions in the coupling constant versus mass space due to their specific sensitivity. Figure 5.1 shows a recent summary of exclusion limits from different experiments assuming only coupling to electron neutrinos $|U_e|^2$. The type I Seesaw mechanism (see section 3.2.2) sets a lower bound in the HNL coupling versus mass plane where the mixing of the HNLs with the active neutrinos becomes too small to produce the observed pattern of neutrino flavour oscillations [147, 148]. The type I Seesaw bound is not visible in fig. 5.1¹, because it is less constraining than the cosmology bound, where the Big Bang Nucleosynthesis is simulated with HNLs present [149]. The bounds on the unitarity of the PMNS matrix come from a global fit to electroweak and flavour observables [150, 151]. The low mass range in fig. 5.1 is dominated by mostly fixed target experiments that study HNLs produced in charged pion decays up to about $m_{\mathcal{N}} \sim 0.14 \text{ GeV}$ [152, 153], charged kaon decays for $0.1 \lesssim m_{\mathcal{N}} \lesssim 0.5 \text{ GeV}$ [154, 155], and charged D and D_s -meson decays for up to $m_{\mathcal{N}} \sim 2 \text{ GeV}$ [156, 157]. In contrast, collider experiments such as DELPHI [158], ATLAS [10, 159], and CMS [160, 161] take over at higher masses in the GeV range up to the electroweak scale for HNLs produced in Z and $W^{\pm}$ boson decays. Similar summaries of the exclusion limits in the HNL coupling versus mass plane for muon-only coupling $|U_{\mu}|^2$ and tau-only coupling $|U_{\tau}|^2$ are given in figs. 5.2 and 5.3. Currently, there are no limits for the tau-only coupling $|U_{\tau}|^2$ in either ATLAS or CMS.

¹The type I Seesaw bound is shown in fig. 8.2 for the total coupling $|U|^2$ as a function of HNL mass together with the expected sensitivities of numerous proposed future experiments.




Figure 5.1: Constraints on $|U_e|^2$ as a function of HNL mass m_N (electron coupling only, Majorana HNLs). Image taken from ref. [163]. The exclusion limits (from left to right) are taken from: The **PIENU** experiment at TRIUMF and **KENU** studying HNLs from pion decays [152, 153], **cosmology** bounds from Big Bang Nucleosynthesis simulations [149], the **Borexino** neutrino experiment [164], the **TRIUMF** limits [165], the **NA62** fixed target experiment at CERN studying HNLs from kaon decays [154], the ND280 detector of the **T2K** experiment studying HNLs from kaon decays [155], the **DELPHI** experiment studying HNLs from Z decays [158], the **Belle** experiment at the KEKB e^+e^- collider [166], the **CHARM** experiment studying HNLs from D meson decays [156], the **BEBC** WA66 beam dump experiment [157], the limits derived in this thesis (**ATLAS2022**, full LHC Run2 data) [10], the previous ATLAS limits that include prompt HNLs from W boson production (**ATLAS2019**, partial LHC Run2 data) [159], the **CMS** experiment (displaced, 2022, full LHC Run2 data) [161] and (prompt, 2018, partial LHC Run2 data) [160], **PMNS unitarity** [150, 151], the **L3** experiment at LEP [167].

The search discussed in this thesis focuses on long-lived HNL signatures. Here, the HNL travels macroscopic distances up to $\mathcal{O}(100)$ mm. Higher lifetimes of the HNL help to probe smaller couplings, as the coupling constant squared scales inversely proportional to the mean lifetime τ_N [162]:

$$|U_\alpha|^2 \propto \frac{1}{c\tau_N \cdot m_N^5} \quad (5.1)$$

Next to the parameterisation in ref. [162], the HNL coupling dependence of m_N^5 becomes clear when comparing it with the three-body muon decay $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$ analogous to the three-body HNL decay $N \rightarrow ll\nu$. In the muon decay, the muon lifetime τ_μ is related to its mass m_μ by $\Gamma(\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu) = 1/\tau_\mu = G_F^2 m_\mu^5 / (192\pi^3)$, where Γ is the decay width and G_F the Fermi constant. G_F is related to the weak coupling constant g_w by the W boson mass m_W : $G_F = \sqrt{2}g_w^2 / (8m_W^2)$. Therefore, we see a similar mass dependency for the muon three-body decay as for the HNL decay in eq. (5.1): $G_F^2 \propto 1/(\tau_\mu m_\mu^5)$ [27].

As shown in eq. (5.1), higher HNL lifetimes lead to smaller couplings, which in turn lead to small cross-sections ($\sigma_N \propto |U_\alpha|^2 \times \text{factor that depends on mass}$).

These constraints lead to bean-shaped exclusion contours as shown in fig. 5.1 for the ATLAS 2022 search [10], which is presented in this thesis. The upper part of the contour is limited by constraints on the decay position in the detector of the HNL to exclude SM decays that happen almost directly at the




Figure 5.2: Constraints on $|U_{\mu}|^2$ as a function of HNL mass m_N (muon coupling only, Majorana HNLs). Image taken from ref. [163]. Exclusion limits (from left to right) are taken from: **cosmology** bounds from Big Bang Nucleosynthesis simulations [149], the **PSI** experiment [168], the **PIENU** experiment [169], the **LSND** neutrino experiment [170], the **T2K** experiment [155, 171], the μ **BooNE** experiment at Fermilab [172], **KEK** [153], the **NA62** experiment [173], the **BEBC** beam dump experiment [174], the **NuTeV** experiment at Fermilab [175], the **NA3** experiment [176], the **DELPHI** experiment [158], the limits derived in this thesis (**ATLAS2022**, full LHC Run2 data) [10], the previous ATLAS limits that include prompt HNLs from W boson production (**ATLAS2019**, partial LHC Run2 data) [159], the **CMS** experiment (displaced, 2022, full LHC Run2 data) [161] and (prompt, 2018, partial LHC Run2 data) [160], **PMNS unitarity** [150, 151].




Figure 5.3: Constraints on $|U_{\tau}|^2$ as a function of HNL mass m_N (tau coupling only, Majorana HNLs). Image taken from ref. [163]. Exclusion limits (from left to right) are taken from: **cosmology** bounds from Big Bang Nucleosynthesis simulations [149], **Atmospheric ν** bounds [177], **Borexino** bounds [178], the **CHARM** experiment [179], the **CHARM** experiment [180], the **ArgoNeuT** experiment [181], the **BEBC** experiment [157], the **BaBar** experiment [182], the **DELPHI** experiment [158], **PMNS unitarity** [150, 151].




Figure 5.4: Feynman diagrams of the leptonic signature considered in the search. The HNL $\mathcal{N}$ is assumed to be produced via W boson decay and then decays itself into two leptons l_β , l_γ and a neutrino ν_γ via an off-shell W^* (charged current). Depending on the model, lepton number conserving (LNC, left) and lepton number violating (LNV, right) decays are considered. Equivalent processes with an initial W^- are valid as well. The crosses $\times$ indicate the mixing between the SM neutrinos and the HNL.

main pp collision point. Constraints on the mass of the displaced vertex limit the contour towards lower masses. Further constraints on the transverse distance of the displaced vertex to the beam line restrict the contour towards large couplings in a functional form of $|U|^2 \propto m_{\mathcal{N}}^{-5}$ for $c\tau_{\mathcal{N}} = \text{const}$. Those selections will be discussed in detail in section 5.6.3. Most interesting are the right ('high' HNL mass) and lower (low couplings) regions of the contour, where we constantly try to push the limits. We are limited by the size of the ATLAS inner tracker when it comes to HNL decays with high displacements with respect to the beam line. Those size limitations mean we lose sensitivity at mean HNL decay lengths of $c\tau_{\mathcal{N}} \approx 100 \dots 1000$ mm. The right side of the contour is limited because of the decrease in coupling strength (that is proportional to the HNL production cross-section) with increasing mass (eq. (5.1)). In contrast to the ATLAS 2022 contour, the CMS 2022 [161] contour shown in fig. 5.1 is not closed at the top. The reason for that is the different analysis strategy that allows for less stringent constraints on the transverse HNL decay position. The CMS strategy, in turn, requires a more complex treatment of SM backgrounds. We in ATLAS did not focus on excluding the 'high' coupling region, as DELPHI already excluded it. Our goal in this search was to push the limits towards lower couplings and higher masses.

5.2 SIGNATURE

The search focuses on the leptonic HNL production via $W^\pm$ bosons produced in the pp collisions as shown in fig. 5.4. Depending on the model taken into account, the decay of the HNL can be lepton number conserving (LNC) or lepton number violating (LNV). If, for example, the HNL is assumed to be a Dirac particle, only LNC decays are considered. On the other hand, if the HNL is assumed to be a Majorana particle, both LNC and LNV decays can happen at the same rate.

The W boson decays close to the interaction point to an SM neutrino ν_α and a charged lepton l_α . Because of the production of the lepton l_α close to the interaction point, it is called the *prompt lepton*. The HNL $\mathcal{N}$ is produced by mixing with the SM neutrino ν_α (indicated by the crosses $\times$ in fig. 5.4) and travels macroscopic distances through the inner tracker due to its long lifetime until it decays leptonically (via the mixing with ν_β) into two opposite-charged leptons $l_\beta^\pm, l_\gamma^\mp$ and a neutrino ν_γ . The production of the two leptons $l_\beta^\pm, l_\gamma^\mp$ away from the interaction point leads to a significant displacement compared to the prompt lepton. They are therefore called *displaced leptons*.

The neutrino ν_γ from the HNL decay can't be detected, but the two displaced leptons can be reconstructed. The fact that the two displaced leptons emerge from the same point in space makes it possible to reconstruct a displaced vertex, which is significantly dislocated with respect to the beamline. The signature of the analysis can, therefore, be reduced to events featuring a prompt lepton and a displaced vertex with opposite-sign leptons as summarised in fig. 5.5.

The signal leptons can have any flavour. However, the sensitivity of the analysis to tau leptons is very low due to their challenging reconstruction, so we focus on electrons and muons for this work. Tau leptons have a short proper decay length of $87 \mu\text{m}$ [183] and typically decay before reaching the innermost pixel layer of the ATLAS detector if they are produced at the primary vertex. Tau leptons decay 35% of the time leptonically ($\tau \rightarrow \nu_\tau \nu_\ell \ell$ with $\ell = e, \mu$) to a muon or electron and two neutrinos, and 65% of them decay hadronically ($\tau \rightarrow \nu_\tau$ hadrons), involving one (72%) or three charged pions (22%). The hadronic part of the tau lepton reconstruction is seeded by jets, and the reconstructed pion tracks are matched to the tau lepton candidate via multivariate discriminant techniques [184]. Using information from the reconstructed tracks and calorimeter energy clusters, a recurring neural network separates the tau lepton candidate from jets initiated by other quarks or gluons [185]. The reconstruction of prompt tau leptons is very involved and usually starts at high transverse momentum values compared to the electron or muon reconstruction ($p_T > 20 \text{ GeV}$) [185, 186]. Another issue is to trigger on tau leptons, as the single tau lepton triggers in ATLAS usually start at $p_T \gtrsim \mathcal{O}(100) \text{ GeV}$, we would need to use single electron or muon triggers instead to cover the leptonic tau lepton decay mode. The tau lepton trigger efficiency using the single electron or muon triggers would then, by construction, be less than 35% because we miss the hadronic tau lepton decays and because of the single electron and muon trigger efficiencies. Another crucial step would be transferring the prompt tau lepton reconstruction scheme to a displaced reconstruction scheme, which has not been followed in this work.




Figure 5.5: Schematic illustration of the signature in the ATLAS detector we define as *signal*.




Figure 5.6: Feynman diagram for the HNL decay via an off-shell Z^* boson (neutral current) for the LNC case. Equivalent processes with an initial W^- are valid as well.

We define the analysis channel $(\alpha - \beta\gamma)$ using the flavours of the prompt (α) and the displaced (β, γ) leptons. The focus on muons and electrons results in eight possible channel combinations, of which two are indistinguishable in data:

1. $\mu - \mu\mu$
2. $\mu - \mu e$ and $\mu - e\mu$ (combined because the order of the displaced leptons cannot be probed in data).
The combination of both channels ($\mu - \mu e + \mu - e\mu$) will be labelled with $\mu - \mu e$ in this work.

3. $\mu - ee$

4. $e - ee$

5. $e - e\mu$ and $e - \mu e$ (combined). The combination of both channels ($e - e\mu + e - \mu e$) will be labelled with $e - e\mu$ in this work.

6. $e - \mu\mu$

The HNL can also decay via an off-shell Z boson (neutral current) for the channels where both displaced leptons have the same flavour $\mu - \mu\mu$, $e - ee$, $\mu - ee$, and $e - \mu\mu$ as shown in fig. 5.6. The combination of analysis channels depends on the specific HNL model and will be discussed in the following section.

5.3 SIGNAL MODELS

Particle physics experiments usually assume a single HNL that couples to one neutrino flavour only, as shown in the summary of the exclusion limits in figs. 5.1 and 5.2. But we need at least two HNLs to explain the observed neutrino mass differences and the baryon asymmetry of the universe (BAU), as discussed in detail in section 3.2. That raises the question of why all those experiments motivating their searches for HNLs with neutrino masses and BAU use an unphysical single-flavour coupling model. Figure 5.7 shows allowed regions




Figure 5.7: HNL coupling ratios and allowed regions according to recent neutrino oscillation data [41]. Model benchmark points used in this work are marked by red crosses. The explicit coupling ratios of those benchmark points are summarised in table 5.1. Here: $U^2 := |U|^2$ and $U_\alpha^2 := |U_\alpha|^2$ as defined in eq. (5.3). Image taken from ref. [40].

according to neutrino data [41] for the normal and inverted neutrino mass hierarchies in the HNL coupling ratio space. Red crosses mark the HNL coupling ratios used in this work. The electron-only mixing case is analogue to fig. 5.1 and refers to the case where the HNL production and decay happens only via electron neutrinos. Similarly, the muon-only mixing allows only coupling to muon neutrinos as in fig. 5.2. We use those two single flavour coupling models that assume only one HNL in the analysis to make it comparable to the exclusion limits set by previous experiments. We call those models 1SFH models, for *one single flavour HNL*.

Figure 5.7 shows clearly that the 1SFH models (muon-only and electron-only benchmark points) disagree with the allowed regions indicated as filled areas according to recent neutrino oscillation data [41]. To tackle this problem, we include the so-called Quasi-Dirac HNL model [40] in our set of tested models, as discussed in section 3.2.3. The Quasi-Dirac HNL model considers two HNLs that are highly degenerate in mass and allows for couplings to more than one neutrino flavour. This way, we can probe the normal hierarchy (NH) and inverted hierarchy (IH) benchmark points shown in fig. 5.7 and summarised in table 5.1. Depending on the benchmark point we probe, we call the model 2QDH(IH) or 2QDH(NH).

Figure 5.8 depicts how we combine the different analysis channels to probe the eight models considered in this work.

1. 1SFH(μ) in Dirac limit (LNC only)

analysis channels: $(\mu - \mu\mu + \mu - \mu e + \mu - ee)$




Figure 5.8: Schematic illustration of the possible channel combinations for the eight models considered in this work. The left part shows the channel combinations for the single flavour mixing scenarios assuming only one HNL and the right part shows the combination used for the Quasi-Dirac pair of HNLs introduced in section 3.2.3. All models are evaluated in the Dirac-limit (LNC only) and the Majorana-limit (50% LNC + 50% LNV).

2. **1SFH(μ) in Majorana limit** (0.5 LNC + 0.5 LNV)

analysis channels: $(\mu - \mu\mu + \mu - \mu e + \mu - ee)$

3. **1SFH(e) in Dirac limit** (LNC only)

analysis channels: $(e - ee + e - e\mu + e - \mu\mu)$

4. **1SFH(e) in Majorana limit** (0.5 LNC + 0.5 LNV)

analysis channels: $(e - ee + e - e\mu + e - \mu\mu)$

5. **2QDH(NH) in Dirac limit** (LNC only)

analysis channels: $(\mu - \mu\mu + \mu - \mu e + \mu - ee + e - ee + e - e\mu + e - \mu\mu)$

6. **2QDH(NH) in Majorana limit** (0.5 LNC + 0.5 LNV)

analysis channels: $(\mu - \mu\mu + \mu - \mu e + \mu - ee + e - ee + e - e\mu + e - \mu\mu)$

7. **2QDH(IH) in Dirac limit** (LNC only)

analysis channels: $(\mu - \mu\mu + \mu - \mu e + \mu - ee + e - ee + e - e\mu + e - \mu\mu)$

8. **2QDH(IH) in Majorana limit** (0.5 LNC + 0.5 LNV)

analysis channels: $(\mu - \mu\mu + \mu - \mu e + \mu - ee + e - ee + e - e\mu + e - \mu\mu)$

Models 5-8 agree with the observed neutrino oscillation data, while models 1-4 are included for historical reasons and comparability to other experiments. The only detail missing in this parameterisation is the dark

matter candidate. For this, we would need to add a third very light HNL to our model in the $m_{\mathcal{N}} \sim \text{keV}$ range [7]. We don't do this here because we are not sensitive to low HNL masses $m_{\mathcal{N}} < \mathcal{O}(2) \text{ GeV}$. In addition, the coupling of the dark matter candidate HNL is very small, as discussed in section 3.2.

To calculate the cross-section for the considered signature (see fig. 5.4), we start from the single flavour Majorana HNL and scale to the other models. The ratio for a particular HNL coupling $|U_\alpha|^2$ to the SM neutrinos with respect to the total coupling $|U_{\text{tot}}|^2$ is defined as:

$$x_\alpha = |U_\alpha|^2 / |U_{\text{tot}}|^2. \quad (5.2)$$

With two HNLs $I = 1, 2$ in the model, the coupling to a particular flavour $|U_\alpha|^2$ follows from the sum over all HNL couplings, while the total coupling is defined as the sum over all HNLs I and flavours α [40] (already defined in section 3.2.3 but listed here again for convenience):

$$|U_\alpha|^2 = \sum_I |U_{\alpha I}|^2 \quad \text{and} \quad |U_{\text{tot}}|^2 = \sum_{\alpha, I} |U_{\alpha I}|^2. \quad (5.3)$$

The branching ratio for the HNL production from a W boson $W \rightarrow l_\alpha \mathcal{N}$ can be deduced from the branching ratio for $W \rightarrow l_\alpha \nu$ [187]:

$$BR(W \rightarrow l_\alpha \mathcal{N}) = BR(W \rightarrow l_\alpha \nu) \cdot |U_\alpha|^2 \left(1 - \frac{m_{\mathcal{N}}^2}{m_W^2}\right)^2 \left(1 + \frac{m_{\mathcal{N}}^2}{2m_W^2}\right) \quad (5.4)$$

Starting with the W production cross-section $\sigma(pp \rightarrow W)$ and using equations 5.4 and 5.2, we can get the cross-section for producing a single flavour Majorana HNL $\sigma_{\mathcal{N}}^M$ with:

$$\begin{aligned} \sigma_{\mathcal{N}}^M &= \sigma(pp \rightarrow W) \cdot BR(W \rightarrow l_\alpha \mathcal{N}) \cdot BR(\mathcal{N} \rightarrow l_\beta l_\gamma \nu) \\ &= \sigma(pp \rightarrow W) \cdot BR(W \rightarrow l_\alpha \nu) \cdot x_\alpha \cdot |U_{\text{tot}}|^2 \cdot \left(1 - \frac{m_{\mathcal{N}}^2}{m_W^2}\right)^2 \left(1 + \frac{m_{\mathcal{N}}^2}{2m_W^2}\right) \cdot BR(\mathcal{N} \rightarrow l_\beta l_\gamma \nu). \end{aligned} \quad (5.5)$$

$|U_{\text{tot}}|^2$ and $BR(\mathcal{N} \rightarrow l_\beta l_\gamma \nu)$ are computed according to ref. [188] with the help of our short-time associates Oleg Ruchayskiy and Jean-Loup Tastet. Figure 5.9 shows the HNL coupling as a function of HNL mass as an example for the 1SFH(e) model. The HNL decay branching ratio $BR(\mathcal{N} \rightarrow l_\beta l_\gamma \nu)$ includes all neutrino flavours and is shown in fig. 5.9 for all models and decay channels. The value for the W production cross-section times BR $\sigma(pp \rightarrow W) \cdot BR(W \rightarrow l_\alpha \nu) = (20.6 \pm 0.6) \text{ nb}$ is measured by ATLAS in ref. [189]. Table 5.1 lists the coupling strength ratios x_α for the different benchmark models. Those values are inserted in eq. (5.5) to take the coupling fraction in the HNL production vertex into account.

Once we want to make the transition from the single Majorana HNL cross-section to the 2QDH model, or once we want to combine the LNC and LNV signatures, we need to apply additional factors to the cross-section in eq. (5.5) [40]. Those factors are summarised in table 5.2. The cross-section is then computed for

each analysis channel separately, taking the particular model into account. The benchmark points presented in fig. 5.7 and summarised in table 5.2 were agreed upon during the 2020 Feebly Interacting Particle Workshop held at CERN [190].

Table 5.1: Coupling ratios x_α for the different benchmarks as shown in fig. 5.7.

benchmark point	$x_e = U_e ^2/ U_{\text{tot}} ^2$	$x_\mu = U_\mu ^2/ U_{\text{tot}} ^2$	$x_\tau = U_\tau ^2/ U_{\text{tot}} ^2$
1SFH(μ)	0	1	0
1SFH(e)	1	0	0
2QDH(IH)	1/3	1/3	1/3
2QDH(NH)	0.06	0.48	0.46

Table 5.2: Prefactors that need to be added to the final cross-section according to [40].

	-1SFH-		-2QDH-	
	Majorana	Dirac	Majorana	Dirac
LNC	$\sigma_{\mathcal{N}}^M$	$2\sigma_{\mathcal{N}}^M$	$2\sigma_{\mathcal{N}}^M$	$4\sigma_{\mathcal{N}}^M$
LNV	$\sigma_{\mathcal{N}}^M$	0	$2\sigma_{\mathcal{N}}^M$	0
Total	$2\sigma_{\mathcal{N}}^M$	$2\sigma_{\mathcal{N}}^M$	$4\sigma_{\mathcal{N}}^M$	$4\sigma_{\mathcal{N}}^M$

5.4 SAMPLES

5.4.1 SIGNAL SAMPLES

The underlying principles of Monte Carlo (MC) simulations have been outlined in section 4.5. Here, we will summarise the used event generators and their parameterisation. The signal samples used in this work are generated using the PYTHIA 8.212 event generator [144] with the A14 tune [191] and the NNPDF2.3LO PDF set [192]. The NNPDF2.3LO PDF set is based on a global dataset that includes LHC data, with the corresponding A14 tune delivering the parameters for parton showering and multi-parton interactions.

Simulated minimum bias events are used to model pileup effects of multiple pp interactions per bunch crossing. Those events are simulated using PYTHIA 8.212 in the A3 tune [193] and the NNPDF2.3LO PDF set.

Particles are propagated through the ATLAS detector using the GEANT4 simulation [145]. The detector response itself is simulated using the ATLAS fast (Atlfast-II) simulation [146]. All MC simulations are hosted in the ATLAS simulation infrastructure [143].

PYTHIA 8.212 does not consider spin correlations because the HNL production and decay are treated separately. We use events created with MADGRAPH5_AMC@NLO 2.9.3 [194] in the HEAVYN model [195, 196] to reweight the signal samples [197]. The reweighting procedure is validated by comparing p_T distributions




Figure 5.9: a) HNL coupling constant as a function of HNL mass for three different mean decay lengths and the 1SFH(e) model. b-e) Branching ratios (BR) for the four models in the Majorana-limit as a function of HNL mass for all considered HNL decay channels. The neutrinos in the branching ratios include all neutrino flavours [188]. The BR for $\mathcal{N} \rightarrow \mu e \nu$ ($\mathcal{N} \rightarrow e \mu \nu$) is zero for the 1SFH(e) (1SFH(μ)) model because only the charged current contributes (fig. 5.4). The BR for $\mathcal{N} \rightarrow \mu \mu \nu$ ($\mathcal{N} \rightarrow ee \nu$) is non-zero for the 1SFH(e) (1SFH(μ)) model because of the small contribution from the neutral current decay (fig. 5.6). The BR differences between same-flavour and different-flavour leptons in the HNL decay stem from interferences between charged and neutral current contributions. All BRs take the relative coupling strengths x_α for the HNL decay into account, which explains the lower BRs for the 2QDH models.

from the PYTHIA 8.212 samples with the samples generated by MADGRAPH5_AMC@NLO 2.9.3. The details about this reweighting procedure are described in section 5.4.5.

To probe the HNL coupling $|U|^2$ versus HNL mass $m_{\mathcal{N}}$ space, we generate signal samples for several HNL masses $m_{\mathcal{N}}$ and proper decay lengths $c\tau_{\mathcal{N}}$ ($c\tau_{\mathcal{N}} \propto 1/|U|^2$, see eq. (5.1)). We probe in total ten different mass points $m_{\mathcal{N}} = [3, 4, 4.5, 5, 7.5, 10, 12.5, 15, 17.5, 20]$ GeV and three different proper decay lengths $c\tau_{\mathcal{N}} = [1, 10, 100]$ mm, which means 30 samples per channel and 180 samples in total for all six analysis channels.

The sample production is divided into three MC simulation campaigns, `mc16a`, `mc16d`, and `mc16e`, which correspond to the years 2015-2016, 2017, and 2018 of pp collision data-taking. For the three campaigns, we generated a total of 50k events. Therefore, the total number of generated signal events is 9 Million.

5.4.2 DATA SAMPLES

This search considers the full LHC Run2 pp collision data recorded with the ATLAS detector in 2015-2018. The data from this period corresponds to an integrated luminosity of $L = 139 \text{ fb}^{-1}$ at a centre of mass energy of $\sqrt{s} = 13 \text{ TeV}$. A standard ATLAS event cleaning procedure is applied, which means events have to:

- Pass the Good Run Lists (GRLs) requirement. The GRLs list run numbers that are considered suitable for physics analyses.
- Contain at least one reconstructed primary vertex (PV) with at least two-tracks with a transverse momentum of at least $p_T > 200 \text{ MeV}$ each.
- If more than one PV is reconstructed, the one with the highest $\sum p_T^2$ is used, where the sum is over the tracks forming the PV.
- Pass an event cleaning that checks for errors during data taking in the detector subsystems.

Before using the dataset, it has to undergo a filtering procedure to reduce the computational effort for large-radius tracking and secondary vertexing. The so-called DRAW filter, as described in section 5.6.2, leaves around 10% of the data for the dedicated long-lived particle reconstruction chain.

5.4.3 WHAT ABOUT DISPLACED TAUS?

Technically, the HNL could also decay into tau leptons for the 2QDH models, as the coupling fraction to the tau lepton sector is nonzero (see table 5.1). To test the sensitivity of the analysis for the reconstruction of displaced tau leptons, we simulate signal samples with $m_{\mathcal{N}} = 10 \text{ GeV}$ in the proper decay lengths of $c\tau_{\mathcal{N}} = 1, 10, 100 \text{ mm}$ for the channels $\mu - \tau\tau$ and $e - \tau\tau$. Since the tau leptons can decay leptonically into an electron or muon and two neutrinos via a W boson, we can test our sensitivity for displaced taus by using




Figure 5.10: Truth W boson kinematics for a $\mu\text{--}\mu\mu$ signal sample are shown with $m_{\mathcal{N}} = 10$ GeV and a lifetime $c\tau_{\mathcal{N}} = 10$ mm. Figures are taken from our internal ref. [16].

the selections designed for the standard analysis channels ($\mu\text{--}\mu\mu$, $\mu\text{--}\mu e$, $\mu\text{--}ee$, $e\text{--}ee$, $e\text{--}e\mu$, and $e\text{--}\mu\mu$) (see section 5.6.3, table 5.4). For all six channels, less than one signal event was observed for the tau test MC samples. Therefore, we conclude that the analysis is not sensitive to displaced tau leptons and set the signal acceptance to zero for the 2QDH models.

5.4.4 SIGNAL KINEMATICS

The kinematic distributions at the 'truth' (generator) level help us to design the signal region selections. Truth level here means that we can look at the distributions directly after the event generation, parton shower, and hadronisation steps before the detector response simulation. This helps us to understand the input signal as generated by the event generator without reconstruction effects.

All relevant kinematic distributions on the truth level for the prompt W boson that decays into $\mathcal{N} + l_\alpha$ are shown in fig. 5.10. The m_W , p_T , η , and ϕ distributions are independent of the HNL mass and lifetime parameterisation.




Figure 5.11: The distributions of the true transverse momentum (p_T) for the three charged leptons involved in the Heavy Neutral Lepton (HNL) decay. The displaced leptons, both leading and sub-leading, are ordered by their p_T . All p_T distributions show three HNL mass points $m_{\mathcal{N}} = 5, 7.5, 10$ GeV with a lifetime $c\tau_{\mathcal{N}} = 10$ mm for the $\mu-\mu\mu$ channel. Figures are taken from our internal note ref. [16].




Figure 5.12: Truth DV mass and tri-lepton mass ($m_{\ell\ell\ell}$) distributions for $\mu-\mu\mu$ signals shown for the three mass points $m_{\mathcal{N}} = 5, 7.5, 10$ GeV with a lifetime $c\tau_{\mathcal{N}} = 10$ mm. Figures are taken from our internal note ref. [16].




Figure 5.13: Truth distribution for the radius of the displaced vertex, $L_{xy} = \sqrt{x^2 + y^2} = r_{DV}$, for a 10 GeV $\mu-\mu\mu$ signal with $c\tau_{\mathcal{N}} = 1, 10, 100$ mm. The dashed red lines indicate material layers in the pixel detector and the first layer of the SCT detector. Figure taken from our internal note ref. [16].

Figure 5.11 shows the true prompt lepton l_{α} transverse momentum p_T as well as the leading and sub-leading p_T distributions from the displaced leptons l_{β} and l_{γ} . It's interesting to note that those distributions do not strongly depend on the generated HNL mass $m_{\mathcal{N}}$. Looking at the lepton p_T distributions, we can also get a good idea about the amount of signal events we lose by the trigger thresholds. A reconstructed muon, for example, needs at least $p_T > 26$ GeV to pass the single muon trigger (see section 5.6.2). In addition, some of the displaced leptons can fire the single muon trigger.

The invariant mass of the displaced vertex m_{DV} and the invariant mass $m_{\ell\ell\ell}$ of all three reconstructed leptons $l_{\alpha}, l_{\beta}, l_{\gamma}$ are shown in fig. 5.12. The displaced vertex mass is always smaller than the generated HNL mass $m_{\mathcal{N}}$ because of the missing neutrino. Therefore, the displaced vertex mass distribution depends on $m_{\mathcal{N}}$. The visible three-lepton mass distribution is independent of the HNL mass and always smaller than the W boson mass. Since all signal samples are populated in a dedicated $m_{\ell\ell\ell}$ mass window, a window selection is well-motivated to reduce possible backgrounds.

Figure 5.13 shows the distribution of the displaced vertex radius r_{DV} (or L_{xy}), which is defined as the transverse distance between the displaced vertex and beam pipe in the xy -plane, for an HNL mass of 10 GeV and three different proper decay lengths. The displaced vertex radius is a two-dimensional projection of the three-dimensional decay position. The distribution of the proper lifetimes t of the generated HNLs is exponential, with the HNL lifetime $\tau_{\mathcal{N}}$ as its mean value. The distance between production and decay vertex in the laboratory frame L_{xyz} depends on the boost of the HNL $L_{xyz} = \gamma\beta ct$, with the ratio of the energy of the HNL to the HNL mass $\gamma = E_{\mathcal{N}}/m_{\mathcal{N}}$ and $\beta = v/c$. Samples with a mean decay length of $c\tau_{\mathcal{N}} = 10$ mm and $c\tau_{\mathcal{N}} = 100$ mm populate the region between $0 < r_{\text{DV}} < 300$ mm with displaced vertices, but the r_{DV} -region is too small to surround most decay vertices with a mean decay length of $c\tau_{\mathcal{N}} = 100$ mm.

The most interesting truth-level distributions are those of the HNL itself. Figure 5.14 shows the calculated HNL mass m_{HNL} , p_T , η , and ϕ distributions of the HNL for three different generated HNL masses $m_{\mathcal{N}} = 5, 7.5, 10$ GeV with a proper decay length of $c\tau_{\mathcal{N}} = 10$ mm. While the p_T , η , and ϕ distributions are more or less independent of the generated HNL mass $m_{\mathcal{N}}$, the calculated HNL mass m_{HNL} peaks nicely at the generated values. The calculated (or later 'reconstructed') HNL mass m_{HNL} is used as the main discriminating variable in this analysis and is determined from the W mass spectrum and the visible decay products l_α, l_β , and l_γ only without the necessity to use the neutrino information ν_γ . This is possible by utilising four-momentum (p^μ) conservation of the HNL production $p_W^\mu = p_\alpha^\mu + p_{\mathcal{N}}^\mu$ and the HNL decay $p_{\mathcal{N}}^\mu = p_\beta^\mu + p_\gamma^\mu + p_\nu^\mu$. By rotating the coordinate system such that the new z -axis points towards the neutrino flight direction, it is possible to obtain a quadratic equation with two solutions for the reconstructed HNL mass m_{HNL} that depend on the W boson mass and the energies and momenta of the visible decay products l_α, l_β , and l_γ . We found that only one of the two solutions for m_{HNL} led to realistic values close to the simulated mass $m_{\mathcal{N}}$. Instead of using the measured W boson pole mass as a fixed value in the m_{HNL} calculation, we use the median W mass from its relativistic Breit-Wigner distribution that is kinematically consistent with the charged-lepton decay products to reduce the cases where we can't find a real solution for the quadratic equation of m_{HNL} . This way, we can calculate the reconstructed HNL mass m_{HNL} in fig. 5.14 without any information on the neutrino ν_γ . Further details on the HNL mass calculation are given in section A.2.

5.4.5 ADDITIONAL WEIGHTS ON SIGNAL SAMPLES

After all signal MC samples have been produced with PYTHIA 8.212, we realised that the spin correlations between the prompt lepton l_α and the first displaced lepton l_β that couples directly to the HNL are not taken into account. The reason for this is that PYTHIA 8.212 treats the HNL production and HNL decay separately in the matrix element calculation $|\mathcal{M}_{\text{Pythia}}|^2 = |\mathcal{M}_{\text{production}}|^2 \times |\mathcal{M}_{\text{decay}}|^2$. The PYTHIA 8.212 event generator had already been used in the previous search for displaced HNLs [159]. Therefore, we decided to use the signal samples again for this search.





Figure 5.14: Truth HNL kinematics for $\mu-\mu\mu$ signals are shown for the three mass points $m_{\mathcal{N}} = 5, 7.5, 10$ GeV with a lifetime $c\tau_{\mathcal{N}} = 10$ mm. Figures are taken from our internal note ref. [16].



HNLs are produced with a favoured helicity state correlated with the prompt lepton due to the chiral nature of the weak interaction. The spin is then conserved during propagation since the HNL is a feebly interacting and electrically neutral particle. Therefore, the spin states of l_α and l_β are correlated.

Because it would have been computationally very expensive to regenerate all signal MC samples using an event generator that calculates a combined matrix element $|\mathcal{M}_{\text{production+decay}}|^2$ (such as MADGRAPH5_AMC@NLO 2.9.3), we use a truth reweighting procedure to compensate the kinematic differences that arise from spin correlations. Six different matrix elements are computed with the help of our short-time associates Oleg Ruchayskiy and Jean-Loup Tastet depending on the charged current (CC, via off-shell W as in fig. 5.4) and neutral current (NC, via off-shell Z as in fig. 5.6) contributions and whether we treat lepton number conserving (LNC) or lepton number violating (LNV) decays: $|\mathcal{M}_{\text{CCNC}}|^2$ (charged current + neutral current, LNC or LNV), $|\mathcal{M}_{\text{NC}}|^2$ (neutral current, LNC or LNV), $|\mathcal{M}_{\text{CC}}|^2$ (charged current, LNC or LNV).

For the analysis channels where we have NC and CC contributions ($\mu-\mu\mu$, $e-ee$, $\mu-ee$, $e-\mu\mu$), the weights are calculated depending on the relative coupling strength x_α in the specific model:

$$w_{\text{CC+NC}} = \frac{x_\alpha |\mathcal{M}_{\text{CCNC}}|^2 + (1 - x_\alpha) |\mathcal{M}_{\text{NC}}|^2}{|\mathcal{M}_{\text{Pythia}}|^2} \quad (5.6)$$

For the remaining two analysis channels with mixed displaced lepton flavours ($\mu-\mu e$, $e-e\mu$), only CC contribute, and the weights are calculated using only the CC matrix element:

$$w_{\text{CC}} = \frac{|\mathcal{M}_{\text{CC}}|^2}{|\mathcal{M}_{\text{Pythia}}|^2} \quad (5.7)$$

Figure 5.15 shows the calculated weights for a ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) sample for one channel where both displaced leptons have the same flavour ($\mu-\mu\mu$) and another channel with mixed displaced lepton flavours ($\mu-\mu e$). The difference in weights between w_{CC} and $w_{\text{CC+NC}}$ is expected because of the differences in contributing diagrams. Finally, we validate the applied truth reweighting scheme using MADGRAPH5_AMC@NLO 2.9.3 in the HEAVYN model [195, 196] that takes the spin correlations into account during the event generation by comparing the p_T distributions of the involved leptons. The validations show agreement between the reference p_T distributions created with MADGRAPH5_AMC@NLO and the reweighted PYTHIA p_T distributions. An example for the p_T distributions that are used for validation is shown in fig. 5.16 for the lepton number conserving decays of the $\mu-\mu\mu$ analysis channel. Similar validation distributions for the $\mu-\mu e$ analysis channel and lepton number violating decays that include only the charged current weights are listed in figs. A.1 to A.3. Further details about the described truth reweighting procedure are given in ref. [17]




Figure 5.16: Truth reweighting validation p_T distributions for the four leptons (l_α , l_β , l_γ , ν_γ) between MADGRAPH5_AMC@NLO and PYTHIA for a $\mu\text{--}\mu\mu$ ($m_N = 10$ GeV, $c\tau_N = 10$ mm) sample. The applied weights are w_{CC+NC} in the 1SFH(μ) model ($x_e = 0$, $x_\mu = 1$, $x_\tau = 0$) for the lepton number conserving decays (LNC). Validation distributions for the $\mu\text{--}\mu e$ channel and lepton number violating decays are listed in figs. A.1 to A.3. Figures are taken from our internal note ref. [16].




Figure 5.17: Schematic illustration of the analysis framework and reconstruction chain. The light pink boxes show the standard reconstruction optimised for prompt objects, resulting in the AOD (analysis object data) format. The DRAW filters are applied to the prompt objects, generating the DRAW_RPVLL format with about 10% of the data. The green boxes show the dedicated long-lived particle workflow featuring large-radius tracking and secondary vertexing algorithms. Further skimming and custom lepton-only vertexing are applied in another derivation step called EXOT29. The calculation of analysis-specific variables and most of the event selection is made at the n-tuple/micro-n-tuple steps. The final step marks the statistical analysis where the exclusion limits are set.

5.5 ANALYSIS FRAMEWORK

Figure 5.17 shows the analysis framework schematically. The dedicated long-lived particle analysis stream is highlighted in green. Large-radius tracking and secondary vertexing algorithms are very expensive in terms of CPU, which makes it necessary to filter the data events before running those algorithms. About 10% of the data events are passed along to the RPVLL stream. A second derivation, called EXOT29, is used to run the customised lepton-only secondary vertexing.

As it can be seen from fig. 5.17, the filters that select HNL-like events are applied to objects from the standard (prompt) reconstruction. By construction, this is a more complex task for signatures with tracks that stem from largely displaced vertices. The so-called DRAW filters are a trade-off between maximising the signal selection efficiency and keeping the data rate manageable, mainly because not only the displaced HNL analysis-specific filters are included in the RPVLL data stream. The detailed filter criteria will be listed later in section 5.6.2.

MC simulation and data start at the DRAW_RPVLL data format, which undergoes the dedicated displaced reconstruction steps of large-radius tracking and secondary vertex finding (introduced in sections 4.4.1 and 4.4.2).




Figure 5.18: Impact parameter $|d_0|$ distributions of the displaced vertex tracks for two signal MC samples ($\mu-\mu\mu$ channel, $c\tau_N = 1$ mm and $c\tau_N = 10$ mm) on truth (MC event generator) level. The standard VSI_LEPTONS selection of $|d_0| \geq 2$ mm is indicated by the horizontal dashed lines. Figure taken from our internal note ref. [16].

The *standard* secondary vertex finding algorithm has not been optimised for the leptonic displaced HNL signature. We, therefore, introduce a customised configuration of the secondary vertexing algorithm that is tailored to the needs of this work. Details are later given in this section (section 5.5.1).

Once the custom secondary vertexing has been run, we can pass on the so-called DAOD_EXOT29 files to our analysis framework that calibrates the physics objects, calculates all required event weights and scale factors, and performs the event selection. At each step of the framework, different files are created. The so-called *Ntuples* still contain a list of events with a multidimensional container structure. In contrast, the *Micro Ntuples* are already reduced to a one-dimensional list of selected displaced vertices. Those *Micro Ntuples* are then passed to the statistical analysis framework, where the exclusion limits are calculated. The technical details about the statistical analysis framework, which is the main part of this work, are summarised in section A.6.

5.5.1 CUSTOM SECONDARY VERTEX RECONSTRUCTION

The standard configuration of the vertexing algorithm (VertexSecondaryInclusive, or VSI) in ATLAS long-lived particle searches is introduced in section 4.4.2. By default, we get two vertexing configurations for free in the DAOD_RPVLL stream, a configuration that uses all tracks in the vertex seeding step, called VSI, and a configuration that uses only lepton tracks in the vertex seeding, called VSI_LEPTONS. Both configurations are not optimised for the low track-multiplicity displaced leptonic HNL signature. Studies done for the LHC Run3 large-radius tracking algorithm to reduce fake vertices gave us some clues on where to start in the optimisation. One main point is to apply additional requirements on the input track collection used by the VSI algorithm in the vertex seeding step to reduce the computational effort on fake vertices.




Figure 5.19: This figure shows the total displaced vertex reconstruction efficiency as a function of displaced vertex radius for three channels $\mu-\mu\mu$, $\mu-\mu e$, $e-ee$. Image taken from our PRL publication ref. [10].

We found that imposing the following selections on the input track collection reduces the same-sign background² by 20%:

- The total number of silicon hits should be ≥ 8 , where silicon hits are the sum of pixel and SCT hits.
- $|z_0| \leq 500$ mm
- The quality of the displaced vertex fit should meet the following requirement: $\chi^2_{\text{ vtx}}/\text{dof} \leq 9$.
- The total number of SCT hits should be ≥ 7 .
- The total number of SCT holes should be ≤ 1 .
- An additional geometrical selection to reduce fake large-radius tracks: $\Delta\Theta \leq 1$ or $|z_0| < 200$ mm. Here, $\Delta\Theta = \Theta - \Theta'$ is the angular difference between the primary Θ and secondary Θ' particle's flight directions, assuming the track is produced in a secondary decay. Further details about this selection are given in ref. [18].

Furthermore, the standard VSI_LEPTONS configuration requires the impact parameter $|d_0|$ to be larger than 2 mm. Looking at the truth track $|d_0|$ of a $c\tau_N = 1$ mm and a $c\tau_N = 10$ mm, $m_N = 10$ GeV, $\mu-\mu\mu$ signal sample, we see that 80% (20%) of the tracks of the 1 mm (10 mm) sample are within $|d_0| \leq 2$ mm. Figure 5.18 shows the distributions of the impact parameter $|d_0|$ on truth-level for the two signal MC samples. Now, reducing the $|d_0|$ selection to zero would improve the signal selection efficiency but also increase the number of background vertices. The optimum appeared to be in the middle, and for the custom vertexing algorithm, we require at least one of the displaced vertex leptons to have $|d_0| > 1$ mm.

²The same-sign background is a data-driven background estimate that uses two-track displaced vertices, where both tracks have the same electric charge, as an estimate for the opposite-sign two-track displaced vertices we expect to see for our signal. The background estimation techniques will be discussed in detail in section 5.7.

The last optimisation step considers the track attachment step of the VSI algorithm. Once seed vertices are formed, tracks that emerge close to the seed vertex can be attached to form a multi-track vertex. The track attachment step is disabled in the standard VSI_LEPTONS configuration. But without the track attachment step, we lose important information about tracks that possibly stem from actual physics events close to the reconstructed vertices. Enabling the track attachment step for our custom vertexing configuration reduced the same-sign background significantly while lowering the signal selection efficiency only slightly. The custom vertexing configuration is called VSI_LEPTONSMOD.

Figure 5.19 shows the total vertex reconstruction efficiency as a function of the displaced vertex radius for the custom vertexing configuration used in this work. The ability to reconstruct vertices is limited by the size of the pixel detector, which causes the efficiency to drop towards zero at $r_{DV} \sim 300$ mm. Similarly, we need some pixel layers to reconstruct secondary vertices successfully. The more pixel layers we have, the better the reconstruction, as the rising slope from small to high r_{DV} values in fig. 5.19 indicates. The reconstruction efficiency starts to increase around $r_{DV} \sim 1$ mm, because we require the impact parameter of at least one of the two displaced leptons to be $|d_0| > 1$ mm. The efficiency peaks around $r_{DV} \sim 40$ mm for all three channels shown here ($\mu-\mu\mu$, $\mu-\mu e$, $e-ee$ channels). Muons have a higher reconstruction efficiency than electrons, which is also reflected in the higher vertex reconstruction efficiency for the $\mu-\mu\mu$ and $\mu-\mu e$ channels than for the $e-ee$ channel.

5.6 EVENT SELECTION

This section summarises the event selection for the different analysis regions utilised in this work. While the signal region (SR) is the most important one to test our signal hypothesis, we use a validation region (VR) to study data-driven background estimation methods and secondary vertexing techniques. Additionally, we define a control region (CR) close to the signal region to constrain the data-driven background estimate. We start by discussing the physics objects that form the signatures we analyse in this work (section 5.6.1). After describing the pre-selection (section 5.6.2), we continue with the signal region selection (section 5.6.3), followed by the signal region selection efficiencies and cutflows (section 5.6.4), and conclude with the specific selection criteria for the control and validation regions (section 5.6.5).

5.6.1 OBJECT DEFINITION

Our analysis includes various physics objects, whose reconstruction is explained in section 4.3: displaced vertices, prompt and displaced muons, and prompt and displaced electrons. Below are their definitions specific to our analysis.

DISPLACED VERTEX

We use a custom algorithm to reconstruct the displaced lepton-only vertices used in this analysis. Only displaced vertices with exactly two-tracks matched to displaced leptons are selected.

PROMPT MUON

We start the selection of the prompt muon with the highest p_T muon. The matched track of the prompt muon needs to emerge close to the main interaction point, which is why we require $|d_0| < 3$ mm and $|z_0 \sin \Theta| < 0.5$ mm. The muon needs to fulfil the combined muon requirements (defined in section 4.3.3) and satisfy the MEDIUM working point (WP) with $p_T > 3$ GeV. The muon track is required not to overlap with any displaced vertex tracks. The event is rejected if the muon candidate overlaps with a prompt electron ($\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} < 0.05$). The DRAW filter requires track-based isolation $p_{T\text{cone}30}/p_T < 0.05$, where p_T is the track transverse momentum. The variable $p_{T\text{cone}30}$ is a measure for track isolation and is formed by summing up the transverse momenta of all tracks within a cone of a given size $\Delta R = 0.3$ around the track.

DISPLACED MUON

We consider a muon displaced if a displaced vertex track is matched to a muon that fulfils the combined muon requirements and that passes the MEDIUM WP with $p_T > 3$ GeV. If the displaced muon track overlaps with an electron candidate and fulfils at least the LOOSE WP, then the track is defined as a muon. The DRAW filter requires track-based isolation $p_{T\text{cone}30}/p_T < 1.0$.

PROMPT ELECTRON

As in the muon case, we start the prompt electron candidate selection with the highest p_T electron that emerges close to the main interaction point ($|d_0| < 3$ mm, $|z_0 \sin \Theta| < 0.5$ mm). The prompt electron candidate needs to have no overlap with any DV tracks and fulfil the LH MEDIUM WP with $p_T > 4.5$ GeV. The event is rejected if the candidate overlaps with a prompt muon ($\Delta R < 0.05$). The DRAW filter requires track-based isolation $p_{T\text{cone}30}/p_T < 0.05$.

DISPLACED ELECTRON

Displaced electron candidates are selected from electrons that are matched to DV tracks. The displaced electron candidates must pass the LH VERY VERY LOOSE WP with $p_T > 4.5$ GeV. The required scale factors used to calibrate the LH VERY VERY LOOSE WP are not provided by the combined performance groups in ATLAS because this is a custom working point. Therefore, scale factors corresponding to the loosest electron WP are applied (LOOSE AND BLAYER with gradient isolation defined in ref. [125]). If the displaced electron track




Figure 5.20: This figure shows two-dimensional histograms of the calibrated displaced electron transverse momentum $p_T^{\text{calibrated-electron}}$ as a function of the transverse momentum of the corresponding track $p_T^{\text{track-only}}$. The p_T distributions show data events without a prompt lepton that contain a displaced electron and a displaced muon with an opposite electric charge that form a two-track displaced vertex (we call this region inverted prompt lepton validation region (VR). The detailed definition follows in section 5.6.5). (a) shows a mismatch between the track and the calibrated electron transverse momentum, which is resolved in (b) by imposing an additional selection of $\Delta p_T/p_T^{\text{track-only}} < 0.5$. Figures are taken from our internal note ref. [16].

overlaps with a muon candidate and fulfils at least the LHLOOSE WP, then the track is defined as an electron. The DRAW filter requires track-based isolation $p_T\text{cone30}/p_T < 1.0$.

We observe discrepancies, shown in fig. 5.20 (a), if we compare the $p_T^{\text{track-only}}$ distributions from inner detector tracks matched to electrons with the calibrated electron $p_T^{\text{calibrated-electron}}$ distributions, which contain the track and the calorimeter information. The track-only variables lead to slightly different results than the calibrated electron-track variables, for example, when calculating the displaced vertex mass. An additional loose selection is imposed on the displaced electrons because of this: $\Delta p_T/p_T^{\text{track-only}} < 0.5$, $\Delta p_T = p_T^{\text{track-only}} - p_T^{\text{calibrated-electron}}$. Figure 5.20 (b) shows the track only and the calibrated electron transverse momentum distributions with the $\Delta p_T/p_T^{\text{track-only}}$ selection applied.

5.6.2 PRESELECTION

TRIGGER

The ATLAS trigger and data acquisition system requires dedicated triggers to select events that contain interesting signatures for physics analyses, as discussed in section 4.2.4. We use single lepton triggers to trigger on the prompt lepton from the W boson decay. Those single lepton triggers are either single electron or single muon triggers, which have a threshold in p_T of about 26 GeV for the LHC Run2 (2015-2018). Important to note here is that also displaced leptons can fire the trigger if their transverse momentum reaches the thresholds of the single lepton triggers. The trigger efficiency is 30 – 40% and is independent of the HNL

mass and lifetime since the prompt lepton p_T spectra depend only slightly on the HNL parameterisation (see section 5.4.4). The single lepton trigger chains we use are named:

- Muons:
 - HLT_MU20_ILOOSE_L1MU15 for the years 2015-2016
 - HLT_MU26_IVARMEDIUM for the years 2015-2018
- Electrons:
 - HLT_E24_LHMEDIUM_L1EM20VH for the years 2015-2016
 - HLT_E26_LHTIGHT_NOD0_IVARLOOSE for the years 2015-2018

DRAW FILTER

The Run2 data must be filtered to reduce the computational load for large-radius tracking and secondary vertexing as described in section 5.5. A combination of four filters formed from the standard physics objects was chosen for this analysis. Those four filters require different combinations of prompt-like and displaced-like muons and electrons³. Table 5.3 summarises the four DRAW filters applied to data to filter for HNL-like signatures. If one or more of the four requirements are met, the event passes the DRAW HNL filter.

Table 5.3: Listing of the DRAW HNL filter criteria that are applied before large-radius tracking and secondary vertexing. The event passes if one of the four criteria is met. All criteria are therefore applied to standard ‘prompt’ objects, hence the differentiation of prompt/displaced-like to prompt/displaced objects defined in section 5.6.1. The modified table is taken from our internal note ref. [16].

		Prompt-like muon displaced-like muon filter	Prompt-like electron displaced-like muon filter	Prompt-like electron displaced-like electron filter	Prompt-like muon displaced-like electron filter
HLT trigger		Single muon triggers	Single electron triggers		Single muon triggers
Prompt-like lepton	p_T	$> 28 \text{ GeV}$			
	$ \eta $	< 2.5			
	Muon type	Combined	-		Combined
	Electron WP	-	LH Loose		-
	Isolation	$\text{ptcone30}/p_T < 0.05$			
Displaced-like lepton	p_T	$> 5 \text{ GeV}$			
	$ \eta $	< 2.5			
	Muon type	Combined, standalone, segment-tagged		-	
	Electron WP	-		None	
	Isolation	$\text{ptcone30}/p_T < 1.0$			
	d_0	$> 0.1 \text{ mm}$ for “good” combined muon “good”: MS-ID matching $\chi^2/\text{dof} < 5$			$> 1 \text{ mm}$

FILTER EMULATOR

Several analyses make use of the DRAW_RPVLL pipeline, which means that many filters are present in the RAW→DRAW_RPVLL step. Unfortunately, the information about the filters passed for each event is not

³The DRAW filter is applied in a stage of data reconstruction where only standard (‘prompt’) objects are available. Therefore, the object definitions of prompt and displaced leptons from section 5.6.1 do not apply here.

stored. In addition, no filter is applied for the signal MC samples. Therefore, we need to emulate the DRAW filter for our signal MC samples such that we don't overestimate our signal yields and apply the same skimming as for data. The difficulty here is that the DRAW filter is applied at a reconstruction step where only standard (prompt) objects are available, while the filter emulator has to rely on overwritten objects that include electrons and muons from large-radius tracking. The resulting mismatch is about 3% and is considered a systematic uncertainty (see section 5.9).

PROMPT LEPTON

Events must contain at least one prompt lepton as defined in section 5.6.1.

DISPLACED VERTEX

Events must include at least one displaced vertex with exactly two-tracks.

TRIGGER MATCHING

At least one muon with MEDIUM quality or one electron with LHMEDIUM quality with $p_T > 27$ GeV needs to be matched to one of the single lepton triggers. Matched here means that we ensure this particular lepton has fired the trigger. This can be either one of the prompt or displaced leptons.

COSMIC MUON VETO

Cosmic muons can be falsely reconstructed as a two-muon displaced vertex with back-to-back tracks. The back-to-back signature makes cosmic muons easy to exclude without sacrificing any signal. To achieve this, we define a separation $S = \sqrt{(\sum \eta)^2 + (\pi - \Delta\phi)^2}$ using the angular track parameters η and ϕ . If both tracks are back-to-back, then $S = 0$. Figure 5.21 shows the separation parameter for $\mu\text{--}\mu\mu$ HNL signal MC samples with a mean decay length of $c\tau_{\mathcal{N}} = 10$ mm and masses of $m_{\mathcal{N}} = 5, 7.5, 10$ GeV. We require $S > 0.05$ to exclude displaced vertices from cosmic muons without affecting the HNL signal distributions.

5.6.3 SIGNAL REGION SELECTION

A summary of all selections that define the final signal region (SR) is given in table 5.4. The SR selection is run for each of the six analysis channels separately. All selection criteria that have not already been described in the preselection (section 5.6.2) are detailed below.

FIDUCIAL VOLUME

We restrict the displaced vertex radius in a fiducial volume between 4 mm and 300 mm: $4 < r_{\text{DV}} < 300$ mm. Starting at 4 mm eliminates all promptly decaying SM backgrounds. Displaced vertices can be reconstructed to the first SCT layer (at $r_{\text{DV}} = 299$ mm). The LRT reconstruction efficiency degrades rapidly beyond




Figure 5.21: Signal MC distributions for the separation between two displaced muons. Back-to-back tracks (cosmic muon background) would have a separation value of zero. The expected number of events for each signal MC sample is shown in the legend. Figures are taken from our internal note ref. [16].

Table 5.4: Description of the selection criteria used in the signal region. The lepton quality selection **LHMEDIUM** for prompt electrons is defined in ref. [130] and **MEDIUM** for muons is defined in ref. [123]. The **DISPLACED** lepton quality requirement for displaced electrons is described in ref. [132]. In the Z mass veto, α and β refer to flavours of the prompt and one of the pair of displaced leptons, respectively. The modified table and caption are taken from our PRL publication ref. [10].

Selection	Value
Event cleaning	Standard ATLAS event cleaning
Trigger	Pass at least one single muon or electron trigger
Trigger matched ℓ	At least one lepton with MEDIUM (or LHMEDIUM) quality
Primary vertex	At least one (standard ATLAS selection)
DRAW filter	Pass any HNL filter
Prompt ℓ quality	MEDIUM (muons) or LHMEDIUM (electrons)
Prompt ℓ impact parameters	$ d_0 < 3 \text{ mm}$ and $ z_0 \sin \theta < 0.5 \text{ mm}$
Cosmic muon veto	$\sqrt{(\Sigma \eta)^2 + (\pi - \Delta \phi)^2} > 0.05$
Displaced lepton-only vertex	At least one
Number of tracks in displaced vertex	2
Fiducial volume	$4 < r_{\text{DV}} < 300 \text{ mm}$
	$ z_{\text{DV}} < 300 \text{ mm}$ (only for ee displaced vertices)
Displaced vertex charge	Opposite-sign (OS) tracks
Prompt + displaced ℓ charge	OS leptons (1SFH in Dirac-limit)
Displaced vertex type	$ee, e\mu$ or $\mu\mu$ vertex
Displaced ℓ quality	MEDIUM (muons), LHVERYVERYLOOSE (electrons)
Material veto	Applied for ee displaced vertices only
Heavy-flavour decay veto	$\mu\mu$ displaced vertices: $m_{\text{DV}} > 5.5 \text{ GeV}$
	ee or $e\mu$ displaced vertices: $m_{\text{DV}} - r_{\text{DV}}$ correlation selection eq. (5.8)
Z mass veto	$m(\ell_\alpha^\pm \ell_\beta^\mp) < 80 \text{ GeV}$ or $m(\ell_\alpha^\pm \ell_\beta^\mp) > 100 \text{ GeV}$, if $\alpha = \beta$
Tri-lepton mass	$40 \text{ GeV} < m_{\ell\ell\ell} < 90 \text{ GeV}$
HNL mass	$m_{\text{HNL}} < 20 \text{ GeV}$

$r_{DV} = 300$ mm because of the requirement to have at least 7 pixel+SCT hits. To maximise the number of good-quality tracks used in vertexing, we limit the displaced vertex radius to 300 mm.

DISPLACED VERTEX CHARGE

The two displaced vertex tracks must have opposite electric charges, as shown in figs. 5.4 and 5.6.

PROMPT + DISPLACED LEPTON CHARGE

This selection further reduces the background while keeping the signal selection efficiency constant. For the Dirac models with single-flavour coupling only (1SFH model), the lepton number is always conserved (LNC only, see fig. 5.4). Therefore, we can require the prompt lepton l_α and the displaced lepton l_β from the HNL decay to have opposite electric charges. The prompt and displaced lepton charge requirement is only applied to the channels $\mu - \mu e$ and $e - e \mu$ because the displaced leptons are experimentally distinguishable.

DISPLACED VERTEX TYPE

In the displaced vertex type selection, we classify the displaced vertex in terms of the lepton flavours of the reconstructed lepton tracks that are matched to the displaced vertex tracks. Depending on the analysis channel, three displaced vertex types are possible: ee , $e\mu$, and $\mu\mu$. Together with the flavour selection of the prompt lepton, the displaced vertex type selection finalises the analysis channel selection.

DISPLACED LEPTON QUALITY

The lepton tracks matched to the displaced vertex tracks must fulfil certain quality criteria, as already introduced in section 5.6.1. For displaced muons, we require the MEDIUM WP. For displaced electrons, we require the custom LHVERYVERYLOOSE WP.

MATERIAL VETO

Hadronic interactions from particles with the detector material and photon conversions $\gamma \rightarrow e^+e^-$ can create many fake e^+e^- displaced vertices close to high-density detector material. For this reason, we apply a three-dimensional material veto map as shown in fig. 5.22 to all ee -type displaced vertices. Displaced vertices close to high material density detector regions shown in fig. 5.22 are removed from the event. The application of the map-based veto significantly suppresses displaced vertices from hadronic interactions at the cost of rejecting approximately 42% of the fiducial volume [198].

HEAVY-FLAVOUR DECAY VETO

In this analysis, we reduce SM backgrounds as much as possible. This is already done by the lower limit of the fiducial volume selection on the displaced vertex radius. Still, heavy-flavour backgrounds such as B-meson




Figure 5.22: Two-dimensional maps of the observed vertex density in regions vetoed by the material map, projected in the (a) $x-y$ plane and (b) $z-R$ plane, with $R := r_{\text{DV}}$. The colour scale is in arbitrary units (a.u.). Figures and caption are taken from ref. [198].

decays can be quite long-lived and require additional requirements on the displaced vertex mass m_{DV} and radius r_{DV} . Depending on the analysis channel, we apply rather strict requirements ($m_{\text{DV}} > 5.5 \text{ GeV}$ for $\mu\mu$ displaced vertices) or more involved requirements ($m_{\text{DV}} - r_{\text{DV}}$ correlation selection for $e\mu$ and ee displaced vertices). The $m_{\text{DV}} - r_{\text{DV}}$ correlation selection for $e\mu$ and ee displaced vertices aims to reduce heavy-flavour background as much as possible while retaining as many signal events as possible. For the correlation selection, we require:

$$\begin{aligned}
 \text{small } r_{\text{DV}}: \quad & m_{\text{DV}} > 5.5 \text{ GeV for } r_{\text{DV}} < \frac{225}{7} \text{ mm} \\
 \text{intermediate } r_{\text{DV}}: \quad & m_{\text{DV}} > 7 \text{ GeV} \cdot \left(1 - \frac{r_{\text{DV}}}{150 \text{ mm}}\right) \\
 \text{large } r_{\text{DV}}: \quad & m_{\text{DV}} > 2 \text{ GeV for } r_{\text{DV}} > \frac{750}{7} \text{ mm}
 \end{aligned} \tag{5.8}$$

Those requirements are motivated by studies conducted in the validation region that will be discussed in section 5.7.1 and indicated as red lines in fig. 5.29.

Z MASS VETO

Sometimes, one lepton from a prompt $Z \rightarrow ll$ decay makes it into a displaced vertex with a third randomly crossing track. Figure 5.23 shows how this looks like for same-sign displaced vertices in data when we plot the invariant mass of the prompt and the opposite-sign displaced lepton. We see distinct mass peaks at the Z mass for the analysis channels where the prompt and one of the displaced leptons have the same flavour. We veto those events by requiring $80 \text{ GeV} < m(l_{\alpha}^{\pm}, l_{\beta}^{\mp}) < 100 \text{ GeV}$ for leptons with the same flavour ($\alpha = \beta$). l_{α} here is the prompt lepton and l_{β} the displaced lepton as in figs. 5.4 and 5.5.




Figure 5.23: The invariant mass between the prompt lepton and same-flavor lepton with opposite-charge in the displaced vertex for same-sign (SS) displaced vertices in data. Applied are the pre-selection requirements, lepton working points, and the heavy-flavour decay veto. We see a distinct peak in the Z mass region for the displayed channels. The numbers in the legend show the observed number of events. Figures are taken from our internal note ref. [16].




Figure 5.24: The event selection efficiency for each mass-lifetime point for two analysis channels. The percentage of the produced MC simulation events that pass all signal region selections is shown. An entry of 0 indicates no events were selected [10]. The selection efficiency of the remaining four analysis channels is shown in fig. A.4. Figures are taken from our PRL publication ref. [10].

TRI-LEPTON MASS

The invariant mass $m_{\ell\ell\ell}$ of the three reconstructed leptons $l_\alpha, l_\beta, l_\gamma$ is always smaller than the W mass because the neutrino escapes undetected. Most of the signal is between 40 GeV and 90 GeV, as the truth tri-lepton mass distribution in fig. 5.12 shows. After reconstruction, the p_T thresholds of the single lepton triggers cut off the lower tail of the $m_{\ell\ell\ell}$ distribution. We require $40 \text{ GeV} < m_{\ell\ell\ell} < 90 \text{ GeV}$ for the signal region.

HNL MASS

The sensitivity drops fast with increasing HNL mass $m_{\mathcal{N}}$ for the displaced HNL signature due to $|U|^2 \propto \frac{1}{m_{\mathcal{N}}^5}$. Therefore, we require the reconstructed HNL mass to be $m_{\text{HNL}} < 20 \text{ GeV}$.

5.6.4 SIGNAL REGION SELECTION EFFICIENCY AND CUTFLOW

Figure 5.24 shows the selection efficiency after applying the full signal region selection to our signal MC samples for the channels $\mu-\mu\mu$ and $\mu-\mu e$. The selection efficiency is the ratio between the number of events that pass the final signal region selection and the total number of generated events per analysis channel. The impact of the heavy-flavour decay veto can be seen on both channels. For the $\mu-\mu\mu$ channel the $m_{\text{DV}} > 5.5 \text{ GeV}$ requirement sets the sensitivity to zero for the $m_{\mathcal{N}} = 3, 4, 4.5, 5 \text{ GeV}$ samples. While the more involved $r_{\text{DV}} - m_{\text{DV}}$ correlation selection retains some sensitivity in the $\mu-\mu e$ channel for low masses and

high lifetimes, while the selection efficiency for $m_{\mathcal{N}} \leq 5$ GeV and $c\tau_{\mathcal{N}} = 1$ mm is zero. In general, the $\mu-\mu\mu$ channel has a higher selection efficiency for high HNL masses and lifetimes due to the higher reconstruction efficiency of $\mu\mu$ compared to $e\mu$ displaced vertices (fig. 5.19) and the higher reconstruction efficiency of muons compared to electrons (figs. 4.16 and 4.19).

Detailed cutflow tables for a ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) signal MC sample scaled to the 1SFH model in the Majorana limit and the full 139 fb^{-1} of LHC Run2 data are given in tables 5.5 and 5.6. We lose a lot of events in the trigger selection due to the high single lepton trigger p_T thresholds, in the DRAW filter step, in vertexing, and in filtering the heavy-flavour background. Even though the mixed channels, such as the $\mu-\mu e$ channel, have a lower selection efficiency than the $\mu-\mu\mu$ channel, we have a higher number of expected events for the $\mu-\mu e$ channel in table 5.5 due to the higher branching ratio in the HNL decay to different flavour leptons compared to same flavour leptons as shown in fig. 5.9. Again, due to the branching ratios, the mixed flavour displaced vertex channels $\mu-\mu e$ and $e-e\mu$ have the highest expected signal yields among the six analysis channels.

As seen from table 5.6, we are left with 0 – 2 observed data events for all channels in the final signal region selection out of $\sim 10^8$ initial events recorded.

Table 5.5: Cutflow for six simulated signal channels showing the weighted number of expected events based on the single-flavour mixing model in the Majorana limit. Each column uses the generated signal sample with the mass hypothesis $m_{\mathcal{N}} = 10$ GeV and proper decay length $c\tau_{\mathcal{N}} = 10$ mm. Table and caption are taken from our PRL publication ref. [10].

Selection	Prompt- μ channels			Prompt- e channels		
	$\mu-\mu\mu$	$\mu-\mu e$	$\mu-ee$	$e-ee$	$e-e\mu$	$e-\mu\mu$
All	60.05	174.44	2.79	60.85	177.42	2.80
Trigger	23.32	64.97	0.99	17.63	54.47	0.88
Primary vertex	23.32	64.97	0.99	17.63	54.47	0.88
DRAW Filter	14.94	39.84	0.55	9.94	32.84	0.57
Medium prompt ℓ	14.06	37.31	0.51	8.11	26.82	0.46
Prompt ℓ overlap	14.03	37.26	0.51	8.08	26.69	0.46
Displaced vertex	5.11	12.05	0.15	2.35	9.15	0.18
Fiducial volume	4.81	11.36	0.14	2.19	8.38	0.17
2-track displaced vertex	4.66	10.87	0.13	2.04	8.07	0.16
OS displaced vertex	4.63	10.77	0.13	2.02	8.01	0.16
Displaced vertex type	4.52	10.51	0.13	2.00	7.86	0.15
Cosmic veto	4.52	10.51	0.13	2.00	7.86	0.15
Lepton p_T	4.39	8.33	0.09	1.35	6.32	0.15
Material veto	–	–	0.06	0.97	–	–
Displaced ℓ quality	4.01	7.48	0.06	0.85	5.69	0.13
Trigger matched ℓ	3.65	6.79	0.05	0.85	5.66	0.13
$m_{\ell\ell\ell}$	3.57	6.55	0.05	0.83	5.50	0.13
Heavy-flavour veto	1.87	3.64	0.03	0.44	2.98	0.07
Z mass veto	1.87	3.64	0.03	0.44	2.96	0.07
m_{HNL}	1.87	3.64	0.03	0.44	2.96	0.07

Table 5.6: Cutflow for the six channels in data showing the number of events passing each successive signal selection for Majorana HNLs. Table and caption are taken from our PRL publication ref. [10].

Selection	Prompt- μ channels			Prompt- e channels		
	$\mu-\mu\mu$	$\mu-\mu e$	$\mu-ee$	$e-ee$	$e-e\mu$	$e-\mu\mu$
All	1.02e+8	1.02e+8	1.02e+8	1.02e+8	1.02e+8	1.02e+8
Trigger	2.20e+7	2.20e+7	2.20e+7	2.19e+7	2.20e+7	2.20e+7
Primary vertex	2.20e+7	2.20e+7	2.20e+7	2.19e+7	2.20e+7	2.20e+7
DRAW Filter	1.63e+7	1.63e+7	1.63e+7	1.62e+7	1.63e+7	1.63e+7
Medium prompt ℓ	8.02e+6	8.02e+6	8.02e+6	6.43e+6	6.45e+6	6.45e+6
Prompt ℓ overlap	6.26e+6	6.26e+6	6.26e+6	4.66e+6	4.68e+6	4.68e+6
Displaced vertex	6.26e+6	6.26e+6	6.26e+6	4.66e+6	4.68e+6	4.68e+6
Fiducial volume	3.21e+6	3.21e+6	3.21e+6	2.44e+6	2.45e+6	2.45e+6
2-track displaced vertex	1.98e+6	1.98e+6	1.98e+6	1.51e+6	1.51e+6	1.51e+6
OS displaced vertex	1.17e+6	1.17e+6	1.17e+6	9.22e+5	9.25e+5	9.25e+5
Displaced vertex type	1.01e+4	9.59e+4	1.01e+6	8.08e+5	6.97e+4	6709
Cosmic veto	1.01e+4	9.59e+4	1.01e+6	8.08e+5	6.97e+4	6706
Lepton p_T	8387	3.68e+4	8.10e+4	6.52e+4	2.69e+4	5652
Material veto	–	–	4.79e+4	3.70e+4	–	–
Displaced ℓ quality	6217	5484	6056	6218	3349	4130
Trigger matched ℓ	5618	4645	5599	6080	3157	3922
$m_{\ell\ell\ell}$	2990	2631	3277	3649	1813	2210
Heavy-flavour veto	16	30	7	15	13	5
Z mass veto	16	29	7	9	12	5
m_{HNL}	2	2	1	2	0	1

5.6.5 VALIDATION REGION AND CONTROL REGION SELECTION

VALIDATION REGION

The validation region used in this work is formed from events without a prompt lepton candidate and, therefore, called inverted prompt lepton validation region. While designing the selection criteria for the signal region, we blinded the opposite-sign (OS) vertices in data to avoid a possible human bias. Once all selection criteria were set and the background was understood, we *unblinded* the signal region and looked into OS displaced vertices in data. To study the behaviour of OS and same-sign (SS) displaced vertices, we need an analysis region similar to the signal region that allows us to look at both OS and SS displaced vertices without unblinding. To achieve this, we look at events with no reconstructed prompt electron or muon that satisfies any lepton-WP or the criteria defining a prompt lepton (see section 5.6.1). We call this requirement the prompt lepton veto. The prompt lepton veto guarantees orthogonality between the validation region and the signal region since HNLs from W boson production always feature a prompt lepton. This way, we can study OS and SS displaced vertices without risking accidentally unblinding and look into data that contains our signal signature. Table 5.7 summarises the selection criteria for the inverted prompt lepton validation region.

However, even with the applied prompt lepton veto, the HNL signal can leak into the validation region. The signal leakage can happen in two ways: first, if the prompt lepton did not get reconstructed, and second, if the HNL is produced via a Z boson instead of a W boson, as demonstrated in fig. 5.25. We checked how many events don't have a reconstructed prompt lepton in our signal MC samples and found that fraction in the final SR selection to be negligible.

To investigate the Z boson production case, we first estimate the number of events we would expect compared to the W boson production $\sigma(pp \rightarrow W) \cdot \text{BR}(W \rightarrow l\nu) = 20.6 \text{ nb}$ [189]. We deduce the cross-section times branching fraction value for the Z production from the ATLAS measurement $\sigma(pp \rightarrow Z) \cdot \text{BR}(Z \rightarrow ll) = 1.98 \text{ nb}$ [189] using $\text{BR}(Z \rightarrow \nu\nu) = 20\%$ and $\text{BR}(Z \rightarrow ll) = 3.7\%$ [183]:

$$\sigma(pp \rightarrow Z) \cdot \text{BR}(Z \rightarrow \nu\nu) = \frac{\sigma(pp \rightarrow Z) \cdot \text{BR}(Z \rightarrow ll)}{\text{BR}(Z \rightarrow ll)} \cdot \text{BR}(Z \rightarrow \nu\nu) = 11.76 \text{ nb} \quad (5.9)$$

Table 5.7: Description of the selection criteria used in the inverted prompt lepton validation region.

Selection	Value
Event cleaning	Standard ATLAS event cleaning
Primary vertex	At least one (standard ATLAS selection)
DRAW filter	Pass any DRAW filter
Prompt lepton veto	0 prompt leptons
Cosmic muon veto	$\sqrt{(\Sigma\eta)^2 + (\pi - \Delta\phi)^2} > 0.05$
Displaced lepton-only vertex	At least one
Number of tracks in DV	2
Fiducial volume	$4 < r_{\text{DV}} < 300 \text{ mm}$
	$ z_{\text{DV}} < 300 \text{ mm}$ (only for ee displaced vertices)
Displaced vertex type	$ee, e\mu$ or $\mu\mu$ vertex




Figure 5.25: Feynman diagram for the HNL production from a Z boson without a prompt lepton that features a two-lepton opposite-sign displaced vertex. Shown here is the LNC process only. The decay via an off-shell Z is possible, too, if the two displaced leptons have the same flavour.

Hence, we expect about half the number of HNLs from Z production than from W production. In addition, events without a prompt lepton would need to pass the DRAW filter, and there is only a subset of filters in the DRAW stream that does not require a prompt lepton trigger. After applying the DRAW filter efficiencies, we are left with only 1-2 signal events in the validation region. Since we have more than 100 events in each channel, we conclude that the signal contamination is negligible.

CONTROL REGION

Additionally to the validation region, we define a control region to study the data-driven background estimate close to the signal region and to constrain the background in the final fit. For the control region, we apply the same selection as for the signal region (table 5.4) except for the last m_{HNL} selection. Here, we require $20 \leq m_{\text{HNL}} \leq 50$ GeV instead. Based on the low sensitivity for high HNL masses, we don't expect any signal contamination in the control region.

5.7 BACKGROUND ESTIMATION

One advantage of long-lived particle searches is that they are free from most SM processes. This section will elaborate on the remaining background sources and how we reject or estimate them. There are several other possibilities to end up with prompt leptons and opposite-sign displaced vertices in the same event. Figure 5.26 summarises the different background sources we consider in this work and classifies them into two categories: backgrounds we can reduce using specific selection criteria and backgrounds we have to estimate. The reducible backgrounds are vertices from material interactions such as $\gamma \rightarrow e^+e^-$, cosmic muons that are falsely reconstructed as a displaced vertex, decays from metastable particles such as B mesons or strange hadrons such as K_s or Λ^0 , and $Z \rightarrow ll$ decays, where one lepton gets falsely paired with a randomly crossing track to a displaced vertex. The cosmic muon background, the background from material interactions, and the background from $Z \rightarrow ll$ decays can be easily vetoed as described in section 5.6.3.




Figure 5.26: Schematic illustration of possible background sources in this work. The background sources are grouped into reducible sources on the left and irreducible sources on the right. The backgrounds on the left are excluded using specific selection criteria while we estimate the remaining background on the right. All backgrounds have the same signature as our signal: an event featuring a prompt lepton (pink) and an opposite-sign displaced vertex (green).

The background source we can not reduce since it is not constrained in displaced vertex mass or radius is called random track crossing background. Here, two-tracks randomly cross somewhere in the inner detector and get falsely reconstructed as a displaced vertex. By chance, a prompt lepton is reconstructed together with the random track crossing vertex. The data-driven method we use to estimate the random track crossing background is described at the end of this section.

To summarise, our background estimation strategy is to define requirements to remove most of the above-mentioned reducible background sources until we are only left with the irreducible sources, namely the random track crossing background in our signal region. Once this is done, we use a data-driven method based on the total number of observed same-sign displaced vertices to predict the expected number of opposite-sign displaced vertices that are formed from randomly crossing tracks. The reason to use same-sign displaced vertices to estimate opposite-sign displaced vertices in data follows from the nature of the random track crossing background. If only vertices formed from randomly crossing tracks are left in the data, we expect to see the same number of opposite-sign displaced vertices N_{OS} and same-sign displaced vertices N_{SS} . $N_{OS} = N_{SS}$ is true for infinite statistics due to the randomness of the creation of the displaced vertices formed by randomly crossing tracks, as tracks with the same electric charge are equally likely to randomly cross and form a displaced vertex as tracks with an opposite electric charge. We validate the data-driven background estimate later in this section.




Figure 5.27: The SM contains a large number of metastable particles. A selection of the SM particle spectrum is shown as a function of mass and proper lifetime. Shaded regions roughly represent the detector-prompt and detector-stable regions of lifetime space for a particle moving close to the speed of light. Figure and caption are taken from ref. [199].

5.7.1 METASTABLE PARTICLE DECAYS

A selection of the SM metastable particle spectrum is shown as a function of mass and proper lifetime in fig. 5.27. Depending on the proper lifetime, some of them have proper decay lengths ($c\tau$) between $1 \mu\text{m}$ and 10 m . Backgrounds from those metastable particle decays that are relevant to this analysis can be highly boosted B mesons that decay semileptonically and form an opposite-sign displaced vertex with two lepton tracks. This double semileptonic decay could be $b \rightarrow cW^*$ and $c \rightarrow s/dW^*$, where the virtual W^* decays further into $W^* \rightarrow l\nu$. Due to their boost, B mesons can be relatively long-lived and enter our signal region in the low m_{DV} mass range between 5 and 6 GeV . J/ψ mesons can be displaced if they are produced by other b hadrons such as B mesons. They would create a narrow peak in displaced vertex mass $m_{\text{DV}} = 3.1 \text{ GeV}$ for same-flavour analysis channels such as $\mu-\mu\mu$, $e-ee$, $e-\mu\mu$, and $\mu-ee$ due to their leptonic decay mode $J/\psi \rightarrow l^+l^-$. For small displaced vertex masses between 1 and 2 GeV , we would also expect backgrounds from strange baryons and τ leptons. But strange baryons almost always decay into hadrons. Below a displaced vertex mass of 1 GeV , K_s and K_L decays dominate.

We use the inverted prompt lepton validation region (section 5.6.5) to study opposite-sign and same-sign displaced vertices in data before unblinding. Figure 5.28 shows two-dimensional histograms for the displaced vertex mass as a function of the displaced vertex radius for the three displaced vertex types $\mu\mu$, $e\mu$, and ee . Most opposite-sign displaced vertices can be found at low values of displaced vertex mass and radius. For $\mu\mu$




Figure 5.28: Two-dimensional histograms of the displaced vertex mass m_{DV} as a function of the displaced vertex radius $L_{xy} := r_{DV}$ of opposite-sign (OS) displaced vertices (DV) in the inverted prompt lepton validation region (VR) for the three different vertex types. In addition to the VR selection criteria defined in section 5.6.5 we require muons to fulfil the MEDIUM WP and electrons the LHVERYVERYLOOSE WP. The material layers of the four pixel and the first SCT layer are indicated as red dashed lines. All figures show data with $L = 139 \text{ fb}^{-1}$. Figures are taken from our internal note ref. [16].

and ee displaced vertices, we see a clear peak from J/ψ decays at $m_{\text{DV}} = 3.1$ GeV. The bulk-like distribution between $m_{\text{DV}} = 0 - 5.5$ GeV is formed from B meson decays. We suspect the band at low m_{DV} values for the ee histogram to stem mainly from photon conversions due to the widespread towards high r_{DV} values.

In general, we expect more opposite-sign than same-sign displaced vertices to be produced from SM processes such as the above-described metastable particle decays. At displaced vertex masses greater than 5.5 GeV, we mostly predict background from random track crossing since none of the previously mentioned long-lived particles are that massive. To test this, we look at the difference between the number of opposite-sign and same-sign displaced vertices in the $m_{\text{DV}} - r_{\text{DV}}$ plane. This difference is shown in fig. 5.29 for $\mu\mu$, $e\mu$, and ee displaced vertices in the inverted prompt lepton validation region. The z -axis is limited to 30 to accentuate minor differences. Indeed, we see an apparent excess of opposite-sign displaced vertices and the main differences at small displaced vertex masses and radii. To reject those opposite-sign displaced vertices produced from the SM metastable particle decays, we define a displaced vertex type specific heavy flavour decay veto, as introduced in section 5.6.3 and indicated by red lines for the different displaced vertex types in fig. 5.29.

For $\mu\mu$ displaced vertices, we define a straight heavy-flavour decay veto, while the correlated $m_{\text{DV}} - r_{\text{DV}}$ veto is mainly motivated by the $e\mu$ displaced vertices but also applied to the ee displaced vertices. In the case of $\mu\mu$ displaced vertices, the correlated $m_{\text{DV}} - r_{\text{DV}}$ veto cannot exclude resonant $J/\psi \rightarrow \mu^+\mu^-$ decays at $m_{\text{DV}} = 3.1$ GeV, which is why the straight veto of $m_{\text{DV}} > 5.5$ GeV is chosen. After removing the backgrounds from metastable particle decays, only the random track crossing background should remain.

5.7.2 RANDOM TRACK CROSSING BACKGROUND

The random track crossing background should appear at equal rates for same-sign and opposite-sign displaced vertices if we have infinite vertices, as mentioned before. So if our assumption is valid and we are only left with the random track crossing background after applying the heavy-flavour decay veto, then the displaced vertex mass distributions should agree for opposite-sign and same-sign displaced vertices.

Figure 5.30 shows the displaced vertex mass distributions after application of the heavy-flavour decay veto for the inverted prompt lepton validation region for the three displaced vertex types. The agreement between the number of opposite-sign and same-sign displaced vertices is overall excellent. An excess remains for $\mu\mu$ opposite-sign displaced vertices at low displaced vertex masses, which will be considered by a systematic uncertainty later on. Now, we can, in principle, directly use the same-sign displaced vertices in data as an estimate for the random track crossing background.

We would be statistically constrained if we only take the same-sign displaced vertices passing the signal region selection (section 5.6.3 modified for same-sign displaced vertices) as a data-driven background estimate for random track crossing. The probability is minimal that two-tracks that randomly cross form a displaced vertex and get matched to muons or electrons while, at the same time, a prompt lepton is present in the




Figure 5.29: Shown here is the difference between the number of opposite-sign (OS) and same-sign (SS) displaced vertices (DV) for the validation region in the displaced vertex mass versus displaced radius plane. The z -scale is limited to 30 to highlight small positive and negative differences. The heavy-flavour decay veto is indicated as a solid red line (defined in eq. (5.8) and section 5.6.3). Figures are taken from our PRL publication ref. [10].




Figure 5.30: The displaced vertex mass distributions in the inverted prompt lepton validation region for opposite-sign (OS) and same-sign (SS) displaced vertices (DV) after applying the heavy flavour decay veto. The distributions show good agreement between OS and SS displaced vertices. A systematic uncertainty considers the remaining disagreements. The data points and their corresponding error bars are slightly shifted for illustrative purposes. Figures are taken from our PRL publication ref. [10].




Figure 5.31: Schematic illustration of the objects we shuffle from different events to re-pair and create new artificial events that form the data-driven estimate for the random track crossing background. We take prompt leptons from events with same-sign displaced vertices with loosened selection criteria (see text). We randomly pair them with opposite-sign displaced vertices from the validation region without a prompt lepton to form the data-driven background estimate.

event. Ideally, we would prefer a continuous background distribution to maximise the benefits of a potential reconstructed HNL mass peak (fig. 5.14).

For this reason, we introduce a prompt lepton shuffling method, assuming that the prompt lepton in the random track crossing background event is uncorrelated with the displaced vertex. The idea is to take prompt leptons from same-sign data events and randomly pair (shuffle) them with opposite-sign displaced vertices from the inverted prompt lepton validation region. This process is illustrated in fig. 5.31 for generating artificial events from the two input objects. We can calculate kinematic variables, including the mass of HNL, using the artificially paired objects. The increase in statistics is limited by the number of prompt leptons and the number of same-sign displaced vertices available for shuffling. To increase the available statistics, we loosen the same-sign displaced vertex selection criteria⁴ to raise the number of available prompt leptons for shuffling. We verified that the kinematic distributions of the looser displaced vertex events remain unchanged upon loosening their selection criteria, thus avoiding potential biases.

Figure 5.32 shows the resulting HNL mass distribution of the prompt-lepton shuffling method (grey squares) for the six analysis channels. Shown in fig. 5.32 are the signal and control regions together and the smoothing effect the shuffling method has on the data-driven same-sign background (grey dots). The shuffled distributions are normalised to the number of observed same-sign events in $0 \leq m_{\text{HNL}} \leq 50$ GeV. Important to note here is that we only observe 0-1 same-sign background events in the signal regions. The observed opposite-sign data events are shown as black dots for completion.

5.7.3 BACKGROUND VALIDATION

Figure 5.30 already shows that the heavy-flavour decay veto is capable of reducing long-lived low-mass SM background to a negligible level since the opposite-sign and same-sign distributions agree. From this agree-

⁴The loosened displaced vertex selection criteria follow the signal-region selections in section 5.6.3 except that: We select same-sign displaced vertices, we do not require a displaced lepton quality, we do not apply the heavy-flavour decay veto, and we impose a loose displaced vertex mass selection of $m_{\text{DV}} > 1$ GeV.




Figure 5.32: HNL mass distributions for the data-driven estimates for the random track crossing background (grey squares), the observed number of same-sign displaced vertex events (grey dots), and the observed number of opposite-sign displaced vertex events (black dots) for the six analysis channels in the signal region (SR, $m_{\text{HNL}} \in [0, 20]$ GeV) and control region (CR, $m_{\text{HNL}} \in [20, 50]$ GeV). The total number of events in both regions is shown behind the legend entries for each plot. Figures are taken from our internal note ref. [16].

Table 5.8: KS test probabilities were calculated for each channel to evaluate the similarity between the shapes of the same-sign background and the distributions obtained by shuffling each same-sign displaced vertex from the modified signal region with a prompt lepton from a different event (fig. 5.33 and fig. 5.34), as explained in the text. The two distributions are considered incompatible if the KS probability is low ($p \ll 1$). However, the results of the KS test show that the original same-sign distributions and the distribution with shuffled prompt leptons are indeed compatible. As a result, pairing a prompt lepton with a same-sign displaced vertex from accidental track crossings is random. The test values are taken from our internal note ref. [16].

Analysis channel	$p_{\text{KS-test}}$ for m_{HNL} (fig. 5.33)	$p_{\text{KS-test}}$ for $m_{\ell\ell\ell}$ (fig. 5.34)
$\mu-\mu\mu$	0.43	0.20
$\mu-\mu e$	0.99	0.23
$\mu-ee$	0.94	0.72
$e-e\mu$	0.80	0.62
$e-ee$	0.99	0.89
$e-\mu\mu$	0.99	0.45

ment, we can conclude that vertices formed from randomly crossing tracks are independent of the charge of those tracks.

Constructing the data-driven random track crossing background estimate with the prompt lepton shuffling method assumes that the random track crossing vertex is not correlated with the prompt lepton in the event.

To validate this assumption, we look at same-sign displaced vertices in the signal region with slightly modified selections: We apply the pre-selection (section 5.6.2) and only apply the displaced lepton quality and material veto selection from section 5.6.3. Additionally, we require $m_{\text{DV}} > 2 \text{ GeV}$. We then compare the reconstructed HNL mass distributions for the original same-sign events, using the original prompt lepton for the HNL mass calculation (grey dots in fig. 5.33) with HNL mass distributions calculated using a random prompt lepton from a different event (coloured triangles in fig. 5.33). The two resulting HNL mass distributions agree well for all six analysis channels. Figure 5.34 shows the same validation scheme for the invariant mass of the three visible leptons $m_{\ell\ell\ell}$.

On top of the visual comparison, we conduct a Kolmogorov-Smirnov (KS) goodness-of-fit test [200] to compare the shape compatibility of the two distributions. The resulting KS probability $p_{\text{KS-test}}$ is listed in table 5.8. The two input distributions are incompatible if the KS probability is much smaller than unity. However, the KS probabilities in table 5.8 indicate that the HNL mass distributions and the three lepton mass distributions are compatible. Therefore, the pairing of a prompt lepton with a vertex from randomly crossing tracks is, in fact, random.

5.8 DISCRIMINATING VARIABLE

This work’s primary discriminating variable between signal and background is the reconstructed HNL mass m_{HNL} . Although the undetected neutrino makes it impossible to reconstruct the full invariant displaced vertex mass directly, we can use our knowledge of the W mass spectrum and the energy-momentum conservation laws to calculate the HNL mass. The details of the HNL mass calculation are given in ref. [10] and for conve-




Figure 5.33: HNL mass distributions of the original same-sign events that serve as a source for the prompt leptons for shuffling (original prompt leptons, grey squares) and the HNL mass distributions from the shuffled events (random prompt lepton, coloured triangles). The shuffling factor f is the ratio of the number of shuffled events to the number of original same-sign background events. The total number of events is shown behind the legend entries for each plot. Figures are taken from our internal note ref. [16].




Figure 5.34: Three-lepton mass distributions of the original same-sign events that serve as a source for the prompt leptons for shuffling (original prompt leptons, grey squares) and the three lepton mass distributions from the shuffled events (random prompt lepton, coloured triangles). The shuffling factor f is the ratio of the number of shuffled events to the number of original same-sign background events. The total number of events is shown behind the legend entries for each plot. Figures are taken from our internal note ref. [16].




Figure 5.35: Shown here are the reconstructed HNL mass distributions in the signal and control regions for the shuffled background estimate, three different HNL signal MC samples with $m_N = 5, 10, 15$ GeV and $c\tau_N = 10$ mm, and the observed opposite-sign data for all six analysis channels. The fine binning here is not used in the final fit, and the shuffled background is normalised by the final two-bin fit described in section 5.10.2. The signal MC samples are scaled to the 2QDH(IH) model in the Majorana limit. Figures are taken from our PRL publication ref. [10].

nience in section A.2. Figure 5.35 shows the reconstructed HNL mass for three different signal MC samples, the data-driven estimate for random track crossing background, and the observed data for all six analysis channels. The signal MC samples leave a nicely peaking signature around their generated HNL mass $m_{\mathcal{N}}$, and it is easy to see that we expect fewer signal events for increasing $m_{\mathcal{N}}$ due to the decrease in HNL coupling squared with increasing HNL mass as shown in eq. (5.1). In the control region, we are only left with the data-driven background estimate and data, which helps us to constrain the background.

We can also observe in fig. 5.35 that the signal region is sparsely populated with the data-driven background in this fine binning. We see this effect even though we increased the statistics using the prompt lepton shuffling method section 5.7.2. Also, looking again at fig. 5.32, we can see that we don't have many same-sign displaced vertices in the signal region to validate the shape of the data-driven shuffled background estimate. Therefore, we decided to use only one bin for the signal region and one bin for the control region in the final fit.

5.9 SYSTEMATIC UNCERTAINTIES

In this section, we will list the different sources of systematic uncertainties that arise from the luminosity measurement, pileup, the DRAW filter discrepancy between data and simulation, tracking, displaced vertexing, triggering, lepton reconstruction, and the cross-section calculation. Those uncertainties are applied to the signal MC samples in the final fit and discussed in section 5.9.1. Table 5.9 summarises the systematic uncertainties in the selection efficiency per analysis channel, which ranges from 8 – 33% in total.

In addition to the systematic uncertainties applied to the signal MC samples, we have sources of uncertainty associated with the data-driven estimation of the random track crossing background. Those are related to the remaining disagreements between the number of same-sign and opposite-sign displaced vertices in the inverted prompt lepton validation region (fig. 5.30) and to the fact that we only have a finite number of prompt leptons for shuffling. The relative uncertainties per analysis channel applied to the data-driven background estimate are listed in table 5.10 and discussed in detail in section 5.9.2. The total uncertainties in the data-driven background estimate range from 2 – 79% for the different analysis channels.

5.9.1 SIGNAL SYSTEMATIC UNCERTAINTIES

TRACK RECONSTRUCTION UNCERTAINTIES

We randomly remove a fraction of tracks from the event to estimate the uncertainties from track reconstruction. The track-dropping probability is provided by a central recommendation derived for tracks with a small impact parameter from the combined performance (CP) groups in ATLAS and depends on the track parameters p_T and η [201]. The tracking uncertainties returned from the CP tool are added in quadrature and are found not to depend strongly on the analysis channel, HNL mass or HNL lifetime. Therefore, we add all signal samples to study the tracking uncertainty. Figure 5.36 shows a cut flow of the applied selections and

Table 5.9: Summary of the systematic uncertainties impacting reconstruction efficiencies. The values in the table refer to the impacts on the signal yields for each channel. Table and caption taken from our PRL publication ref. [10].

Uncertainty Source Channel	Maximum Selection Efficiency Uncertainty [%]					
	$\mu-\mu\mu$	$\mu-\mu e$	$\mu-ee$	$e-ee$	$e-e\mu$	$e-\mu\mu$
Integrated luminosity				1.7		
Pileup				3		
Filter discrepancy				3		
Tracking				3		
Displaced vertex reconstruction	11	21	19	20	28	9
Lepton d_0 extrapolation	5	7	7	7	6	4
Trigger efficiency	< 1	1	< 1	< 1	< 1	< 1
Lepton reconstruction and identification	4	9	12	17	15	2
W cross-section				3		
HNL branching fractions and decay				5		
Total				8 – 33		

Table 5.10: The relative systematic uncertainties in the background estimate are estimated by creating shuffled events from SS displaced vertices or prompt leptons taken from events with different DV channels. The total uncertainty is the quadrature sum of the two contributions. The uncertainties are treated as uncorrelated among the different channels. Table and caption taken from our PRL publication ref. [10].

Channel	Relative Uncertainty (%)		
	DV	Prompt ℓ	Total
$e-ee$	35	< 1	35
$\mu-ee$	27	< 1	27
$e-e\mu$	2	1	2
$\mu-\mu e$	5	5	7
$e-\mu\mu$	70	2	70
$\mu-\mu\mu$	79	3	79




Figure 5.36: Reduction of selection efficiency at the displaced vertex type selection after removing a random collection of tracks following the central performance groups recommendations. The signal MC samples are combined for the analysis channels $\mu-\mu\mu$, $\mu-\mu e$, $e-ee$, $e-\mu\mu$, HNL masses $m_{\mathcal{N}} = 3 - 20$ GeV, and proper decay lengths $c\tau_{\mathcal{N}} = 1 - 100$ mm because the tracking variations do not depend on channel, mass or lifetime. The selection efficiency is reduced by 2.6% at the displaced vertex type selection. Figure taken from our internal note ref. [16].

the reduction of events in % for the samples where some tracks have been removed randomly. The loss in signal selection efficiency is evaluated at the displaced vertex type selection and symmetrised around the nominal value of the expected signal yield for the full collection of tracks. In the symmetrisation, we assume that the random loss in tracks generates the same fluctuation in signal selection efficiency as a random addition of tracks would. It is easier to randomly remove tracks than to randomly add tracks. In fig. 5.36, we see a reduction of 2.6%, which motivates the conservative estimate of 3% for the track reconstruction uncertainty. Figure 5.37 shows the tracking uncertainty resolved in the different analysis channels.

DISPLACED VERTEX RECONSTRUCTION UNCERTAINTY

The displaced vertex reconstruction uncertainty accounts for differences in vertex reconstruction between data and simulation and is the most prominent signal uncertainty. We apply a method similar to ref. [202] to estimate the uncertainty. To study the differences between data and simulation, we look at $K_s^0 \rightarrow \pi^+\pi^-$ decays ($c\tau = 26.84$ mm) from dijet simulations generated with PYTHIA 8.186 and data with the inverted prompt lepton validation region selections applied. The two-pion displaced vertex requires a modified vertexing algorithm identical to VSILEPTONS_{MOD}, except that not only lepton tracks are used in the vertex seeding step. Ideally, we would have used the same vertexing configuration, but there are no suitable leptonic alternatives in data with enough statistics. After applying the prompt lepton veto, we select two-track displaced




Figure 5.37: Shown here is the tracking uncertainty in % for the different analysis channels for HNL masses $m_{\mathcal{N}} = 3 - 20$ GeV, and proper decay lengths $c\tau_{\mathcal{N}} = 1 - 100$ mm combined. Figure taken from our internal note ref. [16].

vertices and require the displaced vertex mass to be around $m_{K_S^0} = 497$ MeV: $487 < m_{\text{DV}} < 507$ MeV. Additionally, we require $|\vec{\alpha}| < 0.01$ rad for further background suppression. We define $\vec{\alpha} = \text{angle}(\vec{r}_{\text{DV}}^T, \vec{p}_{\text{DV}}^T)$ as the angle in the transverse plane between the vector pointing from the primary to the displaced vertex $\vec{r}_{\text{DV}}^T = (x_{\text{DV}}, y_{\text{DV}}) - (x_{\text{PV}}, y_{\text{PV}})$ and the transverse momentum of the displaced vertex $\vec{p}_{\text{DV}}^T$. For $\vec{\alpha}$, a sharp peak close to zero for displaced vertices from $K_S^0 \rightarrow \pi^+ \pi^-$ decays is expected. The differences between data and MC simulation in fig. 5.38 are resolved in displaced vertex radius r_{DV} and transverse momentum p_T . All p_T bins of the data/MC simulation ratio are normalised to the first bin(s) at small r_{DV} depending on the statistics available. In addition, the tracking and vertex reconstruction algorithms are expected to be more robust in the bins close to the interaction point. The normalisation bins are indicated by a red line in fig. 5.38. Deviations from unity in fig. 5.38 are taken as the displaced vertex reconstruction uncertainty (symmetrised around the nominal). The displaced vertex reconstruction uncertainty varies between 0 – 28% depending on the signal mass-lifetime point and analysis channel.

DRAW FILTER UNCERTAINTY

As explained in section 5.6.2, the reconstructed objects that the DRAW filter is applied to in data are no longer available at the reconstruction level because the large-radius tracking algorithm overwrites the track containers. The DRAW filter is applied to data events before large-radius tracking and secondary vertexing are run. To correctly estimate the event yields from simulation, we must emulate the filter for the signal MC samples. The emulation for the signal MC samples is done on the track collection that includes large-radius tracking, which results in a mismatch between the filter emulator and the actual DRAW filter. We apply the filter emulation for both data and MC simulations. Table 5.11 summarises the possible outcome with




Figure 5.38: The dominant signal uncertainty is due to differences in the reconstruction of displaced vertices and tracks between data and MC. This is evaluated by comparing $K_s^0 \rightarrow \pi^+ \pi^-$ event yields in the validation region with MC simulations produced with PYTHIA 8.186 in bins of p_T and r_{DV} . The data/MC ratio is normalised to the bin nearest the interaction point, where the tracking and vertex reconstruction algorithms are expected to be most robust. The summarised difference from 1.0 is applied to each signal vertex as a per-event systematic variation. Figure and caption (modified) taken from our PRL publication ref. [10].

Table 5.11: Definitions of pass-fail combinations with and without large-radius tracking (LRT) enabled. We apply the filter emulation on both the simulation and the data. Therefore, we expect to see true positives in data and true positives and false positives in simulation.

	pass DRAW filter	
	LRT disabled (data-like)	LRT enabled (simulation-like)
true positive	pass	pass
true negative	fail	fail
false positive	fail	pass
false negative	pass	fail

(simulation-like) and without (data-like) large-radius tracking enabled and defines the terms *true positive* (pass with and without large-radius tracking enabled), *true negative* (fail with and without large-radius tracking enabled), *false positive* (fail the large-radius tracking disabled and pass the large-radius tracking enabled case), and *false negative* (pass the large-radius tracking disabled and fail the large-radius tracking enabled case). For data events, we expect to see only true positive events $N_{\text{truepositive}}$, while for simulation, both true positive and false positive $N_{\text{falsepositive}}$ events are possible. To derive the systematic corresponding to the DRAW filter mismatch, we can turn on and off large-radius tracking in our MC samples using:

$$\text{DRAW uncertainty} = \frac{N_{\text{truepositive}}}{N_{\text{truepositive}} + N_{\text{falsepositive}}} \quad (5.10)$$

After applying this method to all mass-lifetime points and analysis channels, we set a conservative estimate of 3% for the DRAW filter systematic.

PILEUP UNCERTAINTY

Weights are applied to the signal sample distributions to match the pileup profile from MC simulations to the pileup profile during data-taking. The uncertainty corresponding to the pileup reweighting method is computed using a central tool that is provided by the combined performance groups in ATLAS [116]. The pileup reweighting uncertainty expresses the remaining differences in the pileup conditions between data and simulations. The overall uncertainty is determined by varying the pileup reweighting factor around the nominal value. The resulting uncertainty is estimated to be 3% for all analysis channels.

LEPTON IDENTIFICATION $|d_0|$ EXTRAPOLATION UNCERTAINTY

Displaced leptons are not covered in the standard (prompt) lepton reconstruction efficiencies due to their high impact parameter $|d_0|$. With increasing $|d_0|$, the lepton reconstruction efficiency decreases. To measure the effect of lepton identification for high $|d_0|$, we study the effect of the lepton quality selection in our signal MC samples (similar to ref. [203]). We divide the number of reconstructed leptons that pass the quality selection by the number of all reconstructed leptons. The ratio cancels the effect on the selection efficiency from other

sources, such as displaced vertexing and tracking. We do this for muons and electrons separately and look at the result binned in $|d_0|$. Figure 5.39 shows the resulting distributions of the impact parameter $|d_0|$ for both muons and electrons. For low $|d_0|$, we are in the regime of the standard lepton reconstruction. Therefore, we normalise the two distributions to the first bin ($|d_0| = 0 - 3$ mm). The difference from unity in the ratio is symmetrised and used as the $|d_0|$ extrapolation uncertainty.

W PRODUCTION CROSS-SECTION UNCERTAINTY

The HNL production cross-section is proportional to the W production cross-section (eq. (5.5)), which we use as an input parameter from the ATLAS measurement in ref. [189]. The ATLAS measurement of the W production cross-section is reported with a 3% systematic uncertainty. We apply the 3% systematic uncertainty to the signal selection efficiency.

INTEGRATED LUMINOSITY UNCERTAINTY

The integrated luminosity corresponding to the years of data-taking from 2015-2018 of $L = 139 \text{ fb}^{-1}$ measured with the LUCID-2 detector [204] has an uncertainty of 1.7% [118].

LEPTON RECONSTRUCTION AND IDENTIFICATION; TRIGGER UNCERTAINTY

We use central tools provided by ATLAS to determine the uncertainties related to the lepton calibration, the detector conditions, and the trigger system. Systematic uncertainties related to the trigger system range up to 1% [205, 206]. Uncertainties pertaining to lepton reconstruction and identification range between 2% and 17% [207, 208].


HNL BRANCHING FRACTIONS AND DECAY

To determine the branching ratio of the decay from the HNL to the two leptons and the neutrino $BR(\mathcal{N} \rightarrow l_\beta l_\gamma \nu)$, we use an updated calculation that features perturbative QCD corrections [188, 209] that are not included in MADGRAPH. The calculation has been done by our short-time associates, Oleg Ruchayskiy and Jean-Loup Tastet. We assign a conservative 5% as an uncertainty to the calculation according to their recommendation.


5.9.2 BACKGROUND SYSTEMATIC UNCERTAINTIES

DISPLACED-VERTEX SYSTEMATIC UNCERTAINTY

After application of the heavy-flavour decay veto, we still observe some differences between the same-sign and opposite-sign displaced vertex distributions (see fig. 5.30), indicating that we significantly reduced the SM background, but did not remove non-random track crossing background entirely. We introduce the displaced-



(a) displaced muons



(b) displaced electrons

Figure 5.39: Shown here is the track d_0 for displaced leptons (muons and electrons) with and without the lepton quality selection applied. The distributions are normalised to the first bin ($0 - 3 \text{ mm}$), and the deviations from unity in the higher $|d_0|$ bins are used as a systematic uncertainty, as shown at the bottom of each plot. Modified figures are taken from our internal note ref. [16].

vertex systematic uncertainty to take the remaining differences between opposite-sign and same-sign displaced vertices into account that arise from the residual non-random track crossing background.

The data-driven background estimate is derived by shuffling opposite-sign displaced vertices from the inverted prompt lepton validation region with prompt leptons from same-sign data events that pass the signal region selection. To measure the differences arising from the total charge of the two tracks of the displaced vertex, we repeat the method with same-sign displaced vertices from the validation region and shuffle them with the same prompt leptons. Figure 5.40 (top) schematically shows the shuffling method to derive the displaced-vertex systematic uncertainty. Both the nominal shuffled background and the displaced-vertex systematic HNL mass distributions are shown in fig. 5.41. The difference between the data-driven shuffled background and the displaced vertex HNL mass distribution is taken as the displaced-vertex systematic uncertainty. The displaced-vertex systematic uncertainty is applied to the data-driven background estimate in the final fit.

Important to note here is that the uncertainty is fully correlated between the signal and control region and that the signal and control region are only one bin each. Figures A.5 and A.6 show the displaced-vertex systematic uncertainty in the final signal and control region binning for completeness. Significant shape differences of the systematic and nominal HNL mass distributions in 5.41 result in more considerable systematic uncertainties. We see the most significant differences between opposite-sign and same-sign displaced vertices for $\mu\mu$ displaced vertices in the low m_{DV} mass range (see fig. 5.30). Consequently, we get the largest displaced-vertex systematic of 70% for the $e-\mu\mu$ channel and 79% for the $\mu-\mu\mu$ channel.

PROMPT-LEPTON SYSTEMATIC UNCERTAINTY

The limit on the number of artificial events for the data-driven background estimate is determined by the number of data events with a prompt lepton and a same-sign displaced vertex that pass the signal region selection. We introduce the prompt-lepton systematic to account for the finite number of prompt leptons available for shuffling. To measure the effect, we repeat the shuffling method with opposite-sign displaced vertices from the inverted prompt lepton validation region and shuffle them with prompt leptons from same-sign events of a different analysis channel. For example, prompt leptons from data events with same-sign displaced vertices of channel A (say $\mu-\mu\mu$) are shuffled with opposite-sign displaced vertices from the validation region of channel B (say $\mu-\mu e$). Figure 5.40 (bottom) illustrates the shuffling process schematically. The resulting HNL mass distribution then shows the effect of using a different set of prompt leptons. Figure 5.41 indicates a good shape agreement in HNL mass between the nominal data-driven background estimate and the prompt-lepton systematic distribution for all analysis channels, resulting in a systematic uncertainty of up to 5%. The prompt-lepton systematic uncertainty is also applied to the background in the final fit and fully correlated between the signal and control region (see figs. A.7 and A.8 for plots that show the prompt-lepton systematic uncertainty in the signal and control region binning used in the final fit).




Figure 5.40: Schematic illustration of the shuffling procedures applied to derive the two systematic uncertainties that arise from the data-driven background estimation. The method to derive the systematic uncertainty that corresponds to the remaining mismatch between opposite-sign and same-sign displaced vertices is shown at the top. The procedure to estimate the uncertainty that takes the limited number of prompt leptons into account is shown at the bottom, where channel A and channel B indicate two different analysis channels (for example $\mu-\mu\mu$ and $\mu-\mu e$).




Figure 5.41: The HNL mass distributions in the signal and control region for the data-driven background estimate (solid squares), the observed same sign displaced vertices (solid circles) together with the HNL mass distribution used to estimate the displaced-vertex systematic (background DV systematic, empty squares), and the distribution to estimate the prompt-lepton systematic (background plep systematic, empty circles) for all six analysis channels. The systematic uncertainty is derived using the differences between the data-driven background and the systematic distributions. The total number of events is shown behind the legend entries for each plot. Figures are taken from our internal note ref. [16].

5.10 STATISTICAL INTERPRETATION

Now we have all the ingredients we need to test the eight different HNL signal hypotheses (section 5.3). We select only events that match the signature characteristic for the leptonic decay of HNLs (sections 5.2 and 5.6.3). We have minimised SM background in the signal region and estimated the remaining random track crossing background using a validated data-driven method (section 5.7). We found a powerful discriminator between signal and background in reconstructing the HNL mass despite the absence of the neutrino in the reconstructed event (section 5.8). Finally, we studied and estimated the various sources of systematic uncertainties that impact the signal selection efficiency and that arise from the estimation of the data-driven background estimate (section 5.9).

The following section will outline the results of this work, starting with a short introduction to the statistical method used, followed by a description of the final maximum likelihood fit.

We use a software framework named TREXFITTER [210] that has been developed to perform statistical data analyses following the CLs method [211, 212] within the ATLAS Collaboration. A PYTHON based wrapper (the DHNLLIMITS framework) needed to be created from scratch around TREXFITTER to perform all relevant tasks and to produce the final results. The technical details of DHNLLIMITS are described in section A.6.

5.10.1 CLS METHOD

"A search implicitly implies a duty to attempt to falsify a new physics hypothesis being tested" [212].

The CLs prescription [211, 212] is the standard method for exclusion and discovery in the LHC experiments. Its usage dates back to the Higgs search at the Large Electron-Positron Collider (LEP) at CERN. The goal of a search for a new particle is to either exclude a signal with high certainty or confirm it with high certainty while minimising the chances of false positives and false negatives. Practically, this is done in terms of a hypothesis test. The baseline hypothesis (H_0) is that all we observe is background, and the alternative hypothesis (H_1) claims that we observe signal and background events. The software framework TREXFITTER uses a global likelihood function⁵ (eq. (5.11)) for all bins and nuisance parameters (one for each systematic uncertainty), where $\mathcal{L}(\mu = 0)$ corresponds to the baseline hypothesis and $\mathcal{L}(\mu = 1)$ to the alternative hypothesis.

$$\mathcal{L}(\vec{n}|\mu, \mu_b, \vec{\Theta}, \vec{\gamma}) = \prod_{i \in \text{bins}} P\left(n_i | \gamma_i \mu S(\vec{\Theta}) + \mu_b B(\vec{\Theta})\right) \cdot \underbrace{\prod_{j \in \text{NP}} G(\Theta_j)}_{\text{systematics}} \cdot \underbrace{\prod_{i \in \text{bins}} P(\gamma_i)}_{\text{MC statistical uncertainty}} \quad (5.11)$$

P and G are Poisson and Gaussian distributions, $\vec{n}$ the number of events for all bins with the components n_i being the number of events for each bin, $\vec{\Theta}$ the vector of nuisance parameters (NP) Θ_j (one for each

⁵The basic concepts of probability density functions, likelihood functions and the maximum likelihood method are outlined in section A.1.

systematic uncertainty), S and B are the signal and background predictions, and μ is the parameter of interest (POI, signal strength). The signal and background predictions form a counting experiment in each bin and are, therefore, described by a Poisson distribution. Each systematic uncertainty is added as a Gaussian term to the likelihood function. The $\vec{\gamma}$ factors, added as Poisson terms, take the statistical uncertainties of the MC simulations into account and scale the non-data-driven part of the prediction (S) in the main Poisson term. One γ_i factor is included in the likelihood function per bin. In the likelihood fit, we introduce an additional free-floating normalisation factor μ_b on the data-driven background estimate per analysis channel besides the signal strength μ .

The one-sided profile likelihood ratio $\lambda(\mu)$ is used in the test statistic q_μ , which is the ratio between the maximised likelihood for a given μ and the maximised likelihood for all μ (conditional likelihood $\mathcal{L}(\mu, \hat{\hat{\theta}}_\mu)$ divided by the unconditional likelihood $\mathcal{L}(\hat{\mu}, \hat{\hat{\theta}}_\mu)$, where $\hat{\hat{\theta}}_\mu$ are the best-fit values of the nuisance parameters for the conditional likelihood and $\hat{\mu}$ and $\hat{\hat{\theta}}_\mu$ are the best-fit values of the signal strength and the nuisance parameters for the unconditional likelihood; Θ_μ is the set of nuisance parameters $\vec{\Theta}$ for a particular μ , including the gamma and normalisation factors). The test statistic is motivated by the Neyman-Pearson Lemma [213], which states that the likelihood ratio is the optimal discriminator in hypothesis testing. We take the negative logarithm of the profile likelihood ratio as our test statistic q_μ :

$$q_\mu = -2 \ln \lambda(\mu) \quad (5.12)$$

with

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{\theta}}_\mu)}{\mathcal{L}(\hat{\mu}, \hat{\hat{\theta}}_\mu)} \quad (5.13)$$

$\hat{\mu}$ is constrained in this work between $[0, 100]$. Wilks Theorem [214] states that $q_\mu = -2 \ln \lambda(\mu)$ follows a χ^2 distribution for a large data sample. In the large-statistics limit, the so-called asymptotic solution allows the analytic calculation of the confidence levels CL_{s+b} and CL_b [215]. We observe only a few events per bin in the final HNL mass binning. Consequently, we must rely on a numerical MC-based method to calculate the test statistic distributions. Using those so-called toy MC experiments, TREXFITTER creates probability density functions for the alternative hypothesis $f(q|\mu)$ and the baseline hypothesis $f(q|\mu = 0)$. This way, we test our HNL signal hypothesis. Depending on the result of the hypothesis test, we can either confirm or exclude the HNL hypotheses up to a certain probability. The corresponding confidence levels are calculated as follows:

$$\text{CL}_{s+b} = \int_{-\infty}^{q_{obs}} f(q|\mu) dq \quad (5.14)$$

$$\text{CL}_b = \int_{-\infty}^{q_{obs}} f(q|0) dq \quad (5.15)$$

Poor compatibility with the alternative hypothesis is indicated by small values of CL_{s+b} , while CL_b very close to one shows insufficient compatibility with the baseline hypothesis. Therefore, the combination of both

Table 5.12: Post-fit values of the normalisation factors μ_b and the signal strength μ used in the fit for the eight different HNL signal models in Dirac-limit (DL) and Majorana-limit (ML) for ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm).

	$\mu_b(\mu-\mu\mu)$	$\mu_b(\mu-\mu e)$	$\mu_b(\mu-ee)$	$\mu_b(e-ee)$	$\mu_b(e-e\mu)$	$\mu_b(e-\mu\mu)$	$\mu(\text{HNL})$
1SFH (μ), DL	$1.92^{+0.64}_{-0.52}$	$2.86^{+1.00}_{-0.81}$	$0.50^{+0.44}_{-0.28}$	-	-	-	$0.00^{+0.16}_{-0.00}$
1SFH (e), DL	-	-	-	$0.67^{+0.39}_{-0.29}$	$0.46^{+0.32}_{-0.22}$	$1.01^{+0.60}_{-0.43}$	$0.00^{+0.30}_{-0.00}$
2QDH (NH), DL	$1.92^{+0.64}_{-0.52}$	$3.00^{+0.84}_{-0.71}$	$0.50^{+0.44}_{-0.28}$	$0.67^{+0.39}_{-0.29}$	$0.39^{+0.20}_{-0.15}$	$1.01^{+0.60}_{-0.43}$	$0.00^{+0.20}_{-0.00}$
2QDH (IH), DL	$1.92^{+0.64}_{-0.52}$	$3.01^{+0.85}_{-0.71}$	$0.50^{+0.44}_{-0.28}$	$0.67^{+0.39}_{-0.29}$	$0.39^{+0.20}_{-0.15}$	$1.01^{+0.60}_{-0.43}$	$0.00^{+0.07}_{-0.00}$
1SFH (μ), ML	$1.92^{+0.64}_{-0.52}$	$3.00^{+0.84}_{-0.71}$	$0.50^{+0.44}_{-0.28}$	-	-	-	$0.00^{+0.33}_{-0.00}$
1SFH (e), ML	-	-	-	$0.67^{+0.39}_{-0.29}$	$0.38^{+0.20}_{-0.15}$	$1.01^{+0.60}_{-0.43}$	$0.00^{+0.24}_{-0.00}$
2QDH (NH), ML	$1.91^{+0.64}_{-0.52}$	$3.00^{+0.84}_{-0.71}$	$0.50^{+0.44}_{-0.28}$	$0.67^{+0.39}_{-0.29}$	$0.39^{+0.20}_{-0.15}$	$1.01^{+0.60}_{-0.43}$	$0.00^{+0.22}_{-0.00}$
2QDH (IH), ML	$1.92^{+0.64}_{-0.52}$	$3.00^{+0.85}_{-0.71}$	$0.50^{+0.44}_{-0.28}$	$0.67^{+0.39}_{-0.29}$	$0.39^{+0.20}_{-0.15}$	$1.01^{+0.60}_{-0.43}$	$0.00^{+0.08}_{-0.00}$

confidence levels determines whether we have a strong exclusion ($\text{CL}_b \approx 0.5$, very small CL_{s+b}) or a strong confirmation ($\text{CL}_b \rightarrow 1$, $\text{CL}_{s+b} \approx 0.5$) of the alternative hypothesis [211]. The CL_s value is defined as the ratio between the signal-plus-background confidence level and the background-only confidence level:

$$\text{CL}_s = \frac{\text{CL}_{s+b}}{\text{CL}_b} \quad (5.16)$$

Normalising the observed confidence level for the alternative hypothesis CL_{s+b} to the background-only confidence level CL_b allows obtaining reasonable exclusion limits on the signal strength even when the observed number of events very low such that the Gaussian approximation per bin does not apply [211, 212].

Finally, the 95% confidence limit on the signal strength can be found by solving the equation:

$$\text{CL}_s(\mu_{\text{Limit}}) = 0.05 \quad (5.17)$$

We exclude a simulated signal sample if the corresponding limit on the signal strength is below one:

$$\mu_{\text{Limit}}(m_{\mathcal{N}}, c\tau_{\mathcal{N}}, \text{model}) < 1 \quad (5.18)$$




Figure 5.42: Pre-fit and post-fit HNL mass histograms of $m_N = 10$ GeV, $c\tau_N = 10$ mm HNL signal samples in the two quasi-Dirac normal hierarchy (2QDH (NH)) HNL model in the Majorana limit (including LNC and LNV decays) and the data-driven background estimate for signal and control regions (SR, CR) and all contributing analysis channels. The signal and background expectation is fitted to the observed data (black points). In the combination of all channels, only a single bin is used for the signal regions ($m_{\text{HNL}} \in [0, 20]$ GeV) and control regions ($m_{\text{HNL}} \in [20, 50]$ GeV). After the maximum likelihood fit (post-fit), the best-fit value for the signal strength is zero, and red dash-dotted lines indicate the pre-fit signal expectations. The black hashed area indicates the combined uncertainties on signal and background expectations. The significance of the deviation between expectation and observed data is shown at the bottom of the histograms [216]. The systematic uncertainties in each channel are fully correlated between the signal and control region. Each channel has a free-floating normalisation factor to scale the post-fit background (summarised in table 5.12 together with the post-fit value of the signal strength).

5.10.2 CHANNEL COMBINATIONS AND MAXIMUM LIKELIHOOD FIT

Figure 5.42 shows an example combination histogram for the HNL mass in the signal $m_{\text{HNL}} \in [0, 20]$ GeV and control $m_{\text{HNL}} \in [20, 50]$ GeV regions for the 2QDH (NH) HNL model ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) in the Majorana limit that takes all six analysis channels into account. The signal region selection has been defined in section 5.6.3, and the control region selection in section 5.6.5. The pre-fit histogram shows the expected HNL signal, the data-driven background estimate, and the total uncertainty as filled areas. The black points are the observed opposite-sign data. We use only one bin for the signal regions and one bin for the control regions due to the low number of expected events. For most bins, the significance Z to observe n events is within $\pm 2\sigma$, given a prediction of $(S + B)$ and the total uncertainty on the prediction. The significance is calculated following the recommendation in ref. [216].

The post-fit histogram in fig. 5.42 shows no significant ($Z < \pm 2\sigma$) deviation between expectation and observed data. Additionally, we can see that the post-fit value for the signal strength μ is zero. A summary of all normalisation factors for all eight signal models is given in table 5.12. Similar histograms are created for each point in $(m_{\mathcal{N}}, c\tau_{\mathcal{N}})$. The best-fit value for the signal strength is zero for all mass-lifetime points for the probed signal models. All final histograms for the eight models and the ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) sample are listed in fig. 5.45 before and in fig. 5.46 after the maximum likelihood fit.

The ranking plot (see fig. 5.43) shows the effect of the 15 most impactful nuisance parameters on the signal strength and corresponds to the maximum likelihood fit in fig. 5.42. Nuisance parameters here are the normalisation factors on the data-driven background estimate for each analysis channel $\mu_b(\alpha - \beta\gamma)$, systematic uncertainties, and gamma factors for the statistical uncertainty for each bin. The list in the ranking plot then starts with the nuisance parameter with the highest impact ($\Delta\mu$) on the post-fit value of the signal strength μ . Those pre-fit and post-fit impacts on the signal strength are shown on the upper axis of fig. 5.43. Four fits are required for each row of the ranking plot, while the specific nuisance parameter Θ is fixed to the following values: 1. post-fit value plus pre-fit uncertainty ($\Theta = \hat{\Theta} + \Delta\Theta$), 2. post-fit value minus pre-fit uncertainty ($\Theta = \hat{\Theta} - \Delta\Theta$), 3. post-fit value plus post-fit uncertainty ($\Theta = \hat{\Theta} + \Delta\hat{\Theta}$), and 4. post-fit value minus post-fit uncertainty ($\Theta = \hat{\Theta} - \Delta\hat{\Theta}$). The impact on the signal strength $\Delta\mu$ is then calculated for each of the four nuisance parameter variations by taking the difference to the nominal fit value for the signal strength μ . In fig. 5.42, the pre-fit impact for the displaced-vertex background systematic uncertainty in the $\mu - \mu\mu$ channel is large compared to the other nuisance parameters, which is why the other impacts are not visible at this scale.

The normalisation factors $\mu_b(\alpha - \beta\gamma)$ and γ factors are, by definition, centred at one in the ranking plot, while systematic uncertainties are centred at zero. The fit can compensate for deviations between expected and observed data yields in the different bins by assigning new central values ($\hat{\Theta}$) to the nuisance parameters, which is called a *pull* ($\hat{\Theta} - \Theta_0$), where Θ_0 are the initial central values for the nuisance parameters. The pull of each nuisance parameter is shown in fig. 5.43 on the lower axis normalised to its pre-fit uncertainty $\Delta\Theta$.

Figure 5.43 shows that the normalisation factors and the displaced vertex uncertainty on the data-driven background estimate have the most considerable impact on the signal strength. Also, the background normalisation factor for the $\mu-\mu\mu$ channel $\mu_b(\mu-\mu\mu)$ has the largest pull to compensate for the pre-fit deviation between background and observed data in the $\mu-\mu\mu$ control region (see fig. 5.42). The anti-correlation between the signal and control region and the size of the displaced vertex uncertainty on the background for the $\mu-\mu\mu$ channel compensates for the resulting deviation in the signal region. Anti-correlation here means that the displaced vertex systematic uncertainty on the background shifts the nominal value *down* by 78.5% in the signal region and *up* by 50.4% in the control region for the $+1\sigma$ variation. This can be seen in fig. 5.44. Hence, the slight pull on the $\mu-\mu\mu$ displaced vertex background uncertainty in the ranking plot (fig. 5.43) and the high pre-fit impact on the signal strength that is so large that other colour bars appear to be zero. Both effects lead to a very good post-fit agreement in the $\mu-\mu\mu$ channel between expectation and observation in fig. 5.42. Those ranking plots have been done for all mass-lifetime points and analysis channels in all eight HNL model interpretations. None of them show unexpected nuisance parameter pulls.




Figure 5.43: Shown here is the ranking plot generated by TRExFITTER [210] that lists the 15 most impactful nuisance parameters Θ ordered by their post-fit impact on the signal strength μ , marked as $\Delta\mu$ on the upper axis, when varying the post-fit values of the nuisance parameters by their post-fit uncertainties $\hat{\Theta} \pm \Delta\hat{\Theta}$. The same variation is done with the pre-fit uncertainties $\hat{\Theta} \pm \Delta\Theta$ to determine the pre-fit impact on μ (see text). The lower axis shows each nuisance parameter pull $\hat{\Theta} - \Theta_0$ from the initial value Θ_0 normalised to its pre-fit uncertainty $\Delta\Theta$. The ranking plot corresponds to the fit shown in fig. 5.42 of a ($m_{\mathcal{N}} = 10 \text{ GeV}$, $c\tau_{\mathcal{N}} = 10 \text{ mm}$) sample in the 2QDH (NH) model. The pre-fit impact of the displaced-vertex background systematic on the signal strength μ is very high compared to the other nuisance parameters, so their impact is not visible in the plot.




Figure 5.44: Single-bin HNL mass histogram of the data-driven shuffled background estimate (section 5.7.2) in the final signal region (one bin, $m_{\text{HNL}} \in [0, 20]$ GeV) and control region binning (one bin, $m_{\text{HNL}} \in [20, 50]$ GeV) for the $\mu-\mu\mu$ analysis channel. The black histogram shows the nominal data-driven background estimate (no shape modification to the nominal value of the background estimate has been applied; therefore, 'original'='modified'=nominal). The red and blue histograms indicate the symmetrised displaced vertex systematic uncertainty for $\pm 1\sigma$ (section 5.9.2). The deviation of the systematic uncertainty in % from the nominal background estimate is shown at the bottom of each plot. The systematic uncertainties in the background estimate are fully correlated between the signal and control region. Similar plots for the remaining analysis channels are listed in figs. A.5 and A.6.




Figure 5.45: Listing of the pre-fit histograms for the eight HNL models in Dirac (LNC decays only) and Majorana-limit (LNC and LNV decays). Detailed description in fig. 5.42.




Figure 5.46: Listing of the post-fit histograms for the eight HNL models in Dirac (LNC decays only) and Majorana-limit (LNC and LNV decays). Detailed description in fig. 5.42.

5.10.3 EXCLUSION LIMITS

Since no excess of events above the background prediction is observed, we go ahead and set 95% confidence limits in the HNL coupling constant versus mass plane. To do this, we must test each signal sample, parameterised by $m_{\mathcal{N}}$ and $c\tau_{\mathcal{N}}$, using the CLs method (section 5.10.1) against the background-only hypothesis. The 95% confidence limit on the signal strength μ_{Limit} can be found by solving eq. (5.17), which in turn requires calculating the confidence levels CL_{s+b} and CL_b for different values of μ (eqs. (5.14) and (5.15)). For each value of μ , both probability density functions $f(q|\mu)$ and $f(q|\mu = 0)$ of the test statistic q_μ need to be created numerically in TRExFITTER using 10k toy MC experiments. We create the function $\text{CL}_s(\mu)$ from 10 different μ values for each $(m_{\mathcal{N}}, c\tau_{\mathcal{N}})$ point to determine μ_{Limit} . Figure 5.47 shows the calculated confidence levels for a $(m_{\mathcal{N}} = 10 \text{ GeV}, c\tau_{\mathcal{N}} = 10 \text{ mm})$ sample in the 2QDH (NH) model in the Majorana limit. In total, we must run 10k times ten times two (=200k) pseudo experiments for each point in $(m_{\mathcal{N}}, c\tau_{\mathcal{N}})$, analysis channel, and signal model to create the probability functions needed for the confidence level calculation. The calculation has been parallelised on several computing nodes as outlined in section A.6 to be done in a reasonable timescale.

The HNL signal samples are generated for three different mean decay lengths $c\tau_{\mathcal{N}}$ and ten different masses $m_{\mathcal{N}}$ to probe the relevant space in the HNL coupling versus the HNL mass plane. Figure 5.48 summarises the grid points that undergo the hypothesis testing in the coupling versus mass plane. For very low mean decay lengths ($c\tau_{\mathcal{N}}=0.1 \text{ mm}$) and very high mean decay lengths ($c\tau_{\mathcal{N}}=1000 \text{ mm}$), we expect the signal acceptance to be zero because of the applied displaced vertex selection, especially the fiducial volume requirement $4 < r_{\text{DV}} < 300 \text{ mm}$, and the very low cross-section at high mean decay lengths ($\sigma \propto |U|^2 \propto 1/c\tau_{\mathcal{N}}$). The $m_{\text{DV}} > 2 \text{ GeV}$ requirement for all channels leads to zero signal acceptance for all samples with $m_{\mathcal{N}} \leq 2 \text{ GeV}$. Therefore, dummy points are added in the coupling versus mass plane where $\mu_{\text{Limit}} = 100$ to reflect the absence of sensitivity (blue squares in fig. 5.48). Green squares in fig. 5.48 show the tested $(m_{\mathcal{N}}, c\tau_{\mathcal{N}})$ points, and red circles mark points that are excluded because their limit on the signal strength is below one.

Now, technically, we get six different limits on the signal strength: One observed limit, one expected limit, two limits for the $\pm 1\sigma$ variation of the expectation, and two limits for the $\pm 2\sigma$ variation of the expectation. To obtain the final exclusion contour, we perform cubic spline interpolations of the limits on the signal strength to find the contours in the HNL coupling versus mass plane where $\mu_{\text{Limit}} = 1$. Figure 5.49 shows an example interpolation for the expected 95% confidence limit on the signal strength for the 2QDH (NH) model in the Majorana-limit. The result of the cubic spline interpolation is colour-coded between $0 < \mu < 5$ to highlight the relevant range around $\mu = 1$. The black line shows the expected 95% confidence limit for $\mu_{\text{Limit}} = 1$. The six cubic spline interpolations are combined in fig. 5.50 to the exclusion limit plot that shows the expected and observed 95% confidence limits for the 2QDH (NH) model in the Majorana-limit. The region within the contour is excluded, and the observed limit agrees within $\pm 1\sigma$ with the expectation. The

exclusion limit contours for the seven other HNL models show very good agreement between observation and expectation as well. Those limit contours are listed in fig. 5.51.

Figure 5.52 compares the observed and expected 95% CL contours of the 1SFH(μ) model in Dirac-limit with the 95% CL contour already set in 2019 by ATLAS using only $L = 32.9 \text{ fb}^{-1}$ of the LHC Run2 dataset [159]. The limit on the total coupling $|U|^2$ set by DELPHI [158] is also shown in fig. 5.52. We see a significant improvement in the observed 95% CL contour set in this work compared to the 2019 ATLAS result in the coupling versus mass region below the DELPHI limit. Differences on the top (low mass, high coupling) of the two ATLAS contours stem from a different background estimation technique in the 2019 search, where a less stringent heavy-flavour decay veto was applied. The number of observed events in the signal region of the 2019 search has been too low to observe any background from metastable SM particle decays above $r_{\text{DV}} > 4 \text{ mm}$ and $m_{\text{DV}} > 4 \text{ GeV}$. The channel-specific heavy-flavour decay veto causes the meander-like shape of the exclusion contours set in this work at low HNL masses and high couplings because of the selection that depends both on r_{DV} and m_{DV} (see section 5.6.3).

All eight observed 95% exclusion limits are summarised in fig. 5.53. The form of the exclusion limit contours is shaped by the lower bound on the displaced vertex mass $m_{\text{DV}} > 2 \text{ GeV}$ that sets the sensitivity to zero for $m_{\mathcal{N}} \leq 2 \text{ GeV}$ and the fiducial volume requirement of the displaced vertex radius $4 < r_{\text{DV}} < 300 \text{ mm}$ that is responsible for a closed contour at high coupling values (low lifetimes) and that limits the sensitivity to very low coupling values (high lifetimes).

It is very challenging to probe the high mass, long lifetime area of the coupling versus mass plane because the production cross-section of the HNL drops with about $1/m_{\mathcal{N}}^5$ for increasing HNL mass and with $1/c\tau_{\mathcal{N}}$ for increasing lifetime (lower coupling values). Due to the limited size of the inner tracking detector in ATLAS, the sensitivity for long-lived HNLs ($c\tau_{\mathcal{N}} \approx 1 \text{ m}$) is low since we only see a small part of the exponential distribution of the particle decay.

The strongest limits are observed for the 2QDH (IH) model, where we have equal contributions from all three lepton flavours ($x_e = 0.33, x_\mu = 0.33, x_\tau = 0.33$), which combines the sensitivities of the muon and the electron sector. Similarly, the sensitivity of the 2QDH (NH) model is mostly driven by the muon sector ($x_e = 0.06, x_\mu = 0.48, x_\tau = 0.46$) since we have no sensitivity to taus. We see an offset of the observed limit contours between the 1SFH and 2QDH models for large couplings because of the different number of HNLs involved that cause an additional factor in the total coupling calculation (section 5.3). Only for the 1SFH models in the Dirac-limit, we can reduce the background by a factor of two without affecting the signal selection efficiency by requiring the prompt and displaced lepton to have opposite-charge (see section 5.6.3), which improves especially the observed 1SFH(μ) contour in the Dirac-limit compared to the Majorana-limit.

A separate study summarised in section A.7 reveals that systematic uncertainties are not the limiting factor in this search. The expected 95% confidence limits have been calculated twice for all eight signal models, once with and once without including the systematic uncertainties in the final maximum likelihood fit. Statistical uncertainties are included in both cases. Adding the systematic uncertainties has very little impact on the




Figure 5.47: Different confidence levels (CL_b : background only, CL_{s+b} : signal plus background, and $CL_s = CL_{s+b}/CL_b$) as a function of signal strength $\mu = \sigma/\sigma_{theory}$ for the 2QDH (NH) model in Majorana-limit. All confidence levels in the final limits are calculated using 10k toy MC experiments. This plot was created with 1k toy MC experiments per μ value for each confidence level CL_b and CL_{s+b} . Shown here is a ($m_N = 10$ GeV, $c\tau_N = 10$ mm) signal MC sample. $CL_s(\mu) = 0.05$ marks the 95% exclusion limit on the signal strength μ . The expected limits on μ and their confidence bands are marked on the bottom of each plot.

expected 95% confidence limits. The low expected and observed data yields in the final fit and the minor impact of the systematic uncertainties lead to the conclusion that the sensitivity is mainly statistically limited.

This analysis is the first-ever collider search for displaced HNLs that reports the exclusion limits in realistic physics models such as the 2QDH model. The differences in observed limits show the importance of a model interpretation that agrees with the claims made in the introduction. Only models that contain at least two HNLs can give rise to the observed neutrino mass hierarchies or the baryon asymmetry in our universe. These observations can't be explained by a single HNL that couples only to one lepton flavour.




Figure 5.48: Coupling constant versus mass grid for the 2QDH (NH) model in the Majorana-limit. Points that are run through the TREXFITTER framework (tested points) are marked as green squares. Points associated with a signal strength $\mu < 1$ are excluded from the coupling constant versus the mass plane and marked by a red circle (expected exclusion). The light blue squares without associated signal MC samples (top row of each plot) correspond to a mean lifetime of 0.1 mm. For those, we assume a signal content of zero because of the fiducial volume requirement $4 < r_{DV} < 300 \text{ mm}$. Similar dummy points are added for 2 GeV (1 mm, 10 mm, and 100 mm lifetime). Here, the acceptance will be zero because of the displaced vertex mass requirement. At lifetimes of 1000 mm, the signal acceptance is zero and dummy points are included for masses of 3, 4, 4.5 and 5 GeV. Each hypothesis test has been performed using 10k toy MC experiments. The mass lifetime points (3 GeV, 1 mm) and (4 GeV, 1 mm) in this particular example do not have any signal acceptance in all six channels and are treated as dummy points.




Figure 5.49: Shown here is the cubic spline interpolation of the associated signal strength values to each tested point in the HNL coupling constant versus mass plane for the 2QDH (NH) model in the Majorana-limit. The annotation on each of the tested mass-lifetime points shows the expected limit on the signal strength μ . The coloured contours show the interpolated expected limit on the signal strength between zero and five to underline the quality of the spline interpolation in the region of interest.




Figure 5.50: The cubic spline interpolation on the limits on the signal strength values for the different mass lifetime points (as in fig. 5.49) is performed six times to obtain the expected and observed exclusion contours in the HNL coupling constant versus mass plane. The region within the contour is excluded, assuming the 2QDH (NH) HNL model in the Majorana limit. We can see that the observed 95% confidence limit (solid black line) agrees within $\pm 1\sigma$ with the expectation (dashed black line).




Figure 5.51: Listing of the observed and expected 95% CL contours for the eight different HNL models in the coupling versus mass plane.






Figure 5.52: Shown here are the observed and expected 95% CL contours for the 1SFH(μ) model in the Dirac-limit ($x_e = 0, x_\mu = 1, x_\tau = 0$) compared to the DELPHI limit [158] on the total coupling $|U|^2$ and the limit previously set by ATLAS in the first displaced HNL search [159]. The regions within the observed ATLAS contours and those above the DELPHI limit are excluded.



(a) Dirac limit



(b) Majorana limit

Figure 5.53: Summary of the observed 95% CL contours for the eight different HNL models in Dirac (LNC decays only) and Majorana-limit (LNC and LNV decays) in the HNL coupling versus mass plane.

6 ONGOING SEARCH FOR LONG-LIVED HEAVY NEUTRAL LEPTONS




Figure 6.1: Title slide of the analysis kickoff meeting in March 2022 that marks the start of the ongoing search for long-lived heavy neutral leptons.

After finalising the search for long-lived heavy neutral leptons (HNLs) presented in chapter 5, we started a new search campaign for long-lived HNLs in March 2022 (fig. 6.1) that is ongoing at the time of writing this thesis. With the lessons learned from the published analysis discussed in chapter 5, we started with new ideas, a new strategy, new team members, and me as one of the two analysis contacts who manage the analysis. Among these new ideas are improvements in large radius tracking, the analysis selection, the addition of Run3 data, a new method for deriving the random track crossing estimate, the addition of the semileptonic HNL decay and SM background MC simulations for the low m_{DV} mass region, and the exploration of novel, MS-only vertexing techniques. In this chapter, I'll outline those ideas and plans without showing results, as the studies are ongoing and preliminary. We aim to have a first publication draft by Summer 2024 in the current analysis timeline.

6.1 IMPROVEMENTS IN LARGE RADIUS TRACKING


In the ongoing analysis, we benefit from substantial improvements in the large radius tracking algorithm (LRT) [217]. Figure 6.2 summarises the reduction in CPU time and disc space per event as a function of the mean number of interactions per bunch crossing μ compared to the previous (legacy) version of LRT. The new LRT algorithm significantly reduces low-quality and fake track candidates early in the reconstruction chain. Also, the maximum $|z_0|$ of the input tracks is reduced from 1500 mm to 500 mm, the maximum track $|\eta|$ is reduced from 5 to 3, and the minimum track p_T is increased from 900 MeV to 1 GeV compared to the previous version of LRT (section 4.4.1).

The previous LRT version was computationally too expensive to be run in the main reconstruction chain. Instead, a so-called DRAW filter must be applied to select signal-like signatures from all data events. In the end, only about 10% of all recorded data events were passed on to the LRT algorithm (see fig. 5.17). The DRAW filter step was responsible for a significant reduction in signal selection efficiency (up to factor two, as shown in table 5.5). In the new software release and thanks to the improvements in LRT, we can pass on all data events to the dedicated displaced reconstruction, as shown schematically in fig. 6.3. LRT is part of the standard ‘prompt’ reconstruction chain, and the modified secondary vertex reconstruction algorithm tailored to the needs of our analysis is part of a dedicated derivation for long-lived particle searches called LLP1.


The analysis framework, as shown in fig. 6.3 (ntuple, micro ntuple production and the statistical analysis), needed considerable changes to run in the new software release due to changes in algorithm handling and data structure.

6.2 A NEW ANALYSIS CHANNEL

The two-body semileptonic channel was not included in the previous analysis. The off-shell $W^{\pm*}$ in the HNL decay can also decay hadronically in $q\bar{q}$. Both quarks can form a bound state, such as a charged pion. The particular two-body semileptonic HNL decay, with the HNL decaying into a charged lepton and a charged pion, is shown in fig. 6.4. In general, all semileptonic HNL decays dominate the branching ratio. At the time of writing, we limit our tests to the two-body semileptonic HNL decays that involve a charged pion, forming the analysis channels $\mu - \mu\pi$ and $e - e\pi$. Once those tests are successful, we would also include the mixed coupling channels $\mu - e\pi$ and $e - \mu\pi$. Figure 6.5 shows the branching ratios for leptonic HNL decays and two-body semileptonic HNL decays as a function of HNL mass in the 2QDH(IH) model. The two-body semileptonic branching ratios are very low for masses above $m_{\mathcal{N}} > 4 - 5$ GeV compared to the leptonic branching ratio, which makes those decay channels primarily engaging in the low HNL mass regime. But this is not necessarily a problem because the low-mass region with $m_{\mathcal{N}} > 2$ GeV and a coupling of about $|U_\mu|^2 \sim 10^{-5}$ is still an uncovered white space in the exclusion limit summary plot (fig. 5.2). In any case, entering this low HNL mass region $2 < m_{\mathcal{N}} < 5$ GeV would require an update to the background



(a) CPU



(b) disc space

Figure 6.2: The (a) processing time taken per event to reconstruct the same data events and (b) event size of the LRT reconstruction output versus average pileup, comparing the legacy LRT (R21) and updated LRT (R24) reconstruction software. The high average pileup regime (open markers) indicates data events from a 2018 machine development run that is not included in the good run lists (GRL). Figures and caption are taken from ref. [217].




Figure 6.3: Schematic illustration of the analysis framework and reconstruction chain. The light pink boxes show the standard reconstruction optimised for prompt objects, resulting in the AOD (analysis object data) format. In the new software release (R24), large radius tracking is already included in the standard reconstruction chain. A derivation step reconstructs the required secondary vertices, resulting in the DAOD_LLPI format. The green boxes show the dedicated long-lived particle workflow. The calculation of analysis-specific variables and most of the event selection is made at the n-tuple/micro-n-tuple steps. The final step marks the statistical analysis where the final histograms will be produced, exclusion limits set, or the discovery made. The differences between the software releases can be seen by comparing this schematic with fig. 5.17.




Figure 6.4: Feynman diagrams of the two-body semileptonic signature currently tested in the ongoing search. The HNL $\mathcal{N}$ is assumed to be produced via W boson decay and then decays itself into a lepton l_β and a charged pion $\pi^\pm$ via an off-shell W^* (charged current). Depending on the model, lepton number conserving (LNC, left) and lepton number violating (LNV, right) decays are considered. Equivalent processes with an initial W^- are valid as well.




Figure 6.5: Branching ratios (BR) for the 2QDH(NH) ($x_e = 0.06, x_\mu = 0.48, x_\tau = 0.46$) model as a function of HNL mass for the leptonic (a) and two-body semileptonic decay channels (b). The neutrinos in the branching ratios include all neutrino flavours [188].

estimation and reduction strategy of the previous analysis, where the heavy flavour veto excluded this region entirely or parts of it, dependent on the analysis channel. We aim to estimate the expected data yields in the low-mass region using MC simulations of possible SM background processes and a low-mass extension of the data-driven background estimate for random track crossing background. Now, one could argue that the high expected SM backgrounds in this region and the low branching ratios would diminish the significance of the signal and make the two-body semileptonic channel not worthwhile to include. But the two-body semileptonic channel holds one last trump: the expected sharp peak in the invariant mass of the displaced vertex in signal without the need to reconstruct the HNL mass via the W boson spectra as it was done in the previous analysis. CMS already has public results of a semileptonic analysis using a different strategy, where the two quarks form one or two jets instead of a bound state but did not reach the desired white space region in the mass-coupling space [218]. We do not have sensitivity estimates yet for the two-body semileptonic channel.

6.3 DATA-DRIVEN BACKGROUND ESTIMATE

The efforts of improving the data-driven random track crossing estimate are twofold. First, we rotate tracks in an event to create artificially displaced vertices that would fit our selection criteria. Second, we re-apply the object shuffling method from the previous analysis (section 5.7.2) using the artificially displaced vertices from the track rotation method and prompt leptons from same-sign events as input. The track rotation and artificial vertex finding are done on the analysis level using the VSI algorithm without the need to rerun the derivation step. The validation of the method is ongoing.

6.4 MUON SPECTROMETER VERTEXING

Finally, improving the reconstruction efficiency for long-lived particles requires out-of-the-box thinking. CMS (see ref. [219]) has done that in a very literal way by simply looking at energy deposits in the muon spectrometers without doing any secondary vertexing and got very close to the desired white space region ($m_N > 2 \text{ GeV}$, $|U_\mu|^2 \sim 10^{-5}$). The energy deposits in the CMS muon spectrometers are caused by hadronic and electromagnetic showers from the decay products of HNLs that decay within the muon system, interacting with the steel shielding material of the muon detectors. We can not do the same in ATLAS because we do not have the iron plates in front of the muon chambers as CMS does, which could cause the required showers from the HNL decay products. Instead, we follow a different strategy as illustrated in fig. 6.6. A dedicated MS vertex algorithm would extend the MS tracklets and try to perform a vertex fit. The MS vertexing technique has already been used in the ATLAS search in ref. [220] and is currently being adapted to the HNL analysis needs. Those MS vertices could be outside the inner tracker and develop the sensitivity to displaced vertices beyond the usual limit of the first SCT layer ($r_{DV} \sim 300 \text{ mm}$). Extending the reconstruction efficiency to higher HNL lifetimes would push the sensitivity to lower coupling values.





Figure 6.6: Schematic illustration of the idea to reconstruct displaced vertices using only information from the muon spectrometer (MS). A possible bending of the tracks due to the magnetic field in the ATLAS detector is neglected in this sketch.


6.5 DATA AND MC SIMULATION SAMPLES

Adding LHC Run3 data that, by the time of writing, involves the years of data taking 2022 and 2023 would add an integrated luminosity of 66 fb^{-1} at $\sqrt{s} = 13.6 \text{ TeV}$ to the available statistics. Figure 6.7 shows the mean number of interactions per bunch crossings for LHC Run 1-3 and the cumulative luminosity as a function of time. In Run3, the pileup conditions (interactions that happen in the same bunch crossing) are harsher than they have been for Run2. The increase in the luminosity of Run2 + Run3 would mean an increase of expected signal events by about 1.5 compared to Run2 only.

The Run2 and Run3 leptonic samples are generated with MADGRAPH5_AMC@NLO 3.3.1 in the HEAVYN model [195–197]. The PYTHIA 8.307 event generator [144] with the A14 tune [191] and the NNPDF2.3LO PDF set [192] is used for showering and the particles are propagated through the ATLAS detector using GEANT4 [145]. Using MADGRAPH5_AMC@NLO directly from the beginning avoids the additional spin-correlation reweighting (section 5.4.5). Compared to the previous analysis, we extend the mass-proper decay length grid we probe in signal to $m_{\mathcal{N}} = [1, 2, 3, 4, 5, 7.5, 10, 12.5, 15, 17.5, 20] \text{ GeV}$ and $c\tau_{\mathcal{N}} = [0.1, 1, 10, 100, 1000] \text{ mm}$. For the semileptonic samples, we are still in the testing and validation phase. At the time of writing, we use MADGRAPH5_AMC@NLO 3.3.1 with the matrix element calculated in leading order in the HEAVYN model to generate $\mu - \mu\pi$ and $e - e\pi$ test samples with $m_{\mathcal{N}} = [2, 3] \text{ GeV}$ and $c\tau_{\mathcal{N}} = [100, 1000] \text{ mm}$.



(a) mean number of interactions/bunch crossing



(b) cumulative luminosity Run3

Figure 6.7: a) Luminosity-weighted distributions for the LHC Run 1-3 of the mean number of interactions per bunch crossing. b) Cumulative luminosity versus time delivered to (green) and recorded by ATLAS (yellow) during stable beams for pp collisions at 13.6 TeV centre-of-mass energy in LHC Run 3 (2022-2023). Figures and caption are taken from ref. [221].

PART III

FINAL WORDS

7 SUMMARY

The search campaign for hypothetical, long-lived, right-handed Majorana neutrinos, or heavy neutral leptons (HNLs), in the ATLAS experiment at the Large Hadron Collider (LHC) at CERN is presented. HNLs with masses below the electroweak scale can give rise to neutrino masses, the observed abundance of matter over antimatter after the Big Bang, and they can serve as dark matter candidates.

In the published search for long-lived HNLs, we analysed the full LHC Run2 proton-proton collision dataset recorded between 2015 and 2018 with a centre-of-mass energy of 13 TeV and an integrated luminosity of 139 fb^{-1} . No signal is observed, and exclusion limits are set on the squared HNL coupling values with the left-handed neutrino states for the HNL mass range of between 3 GeV and 15 GeV. The observed exclusion limits in the HNL coupling versus mass plane agree within $\pm 2\sigma$ with the expected exclusion limits, which confirms the correctness of the number of expected events.

The HNLs are produced in W boson decays with a charged lepton and decay further into two charged leptons and a neutrino. Because of their small coupling, the HNLs are long-lived and dedicated reconstruction algorithms are needed to reconstruct tracks and secondary vertices that are significantly displaced with respect to the beamline. The objects reconstructed in each event are the charged lepton from the W boson decay and the two oppositely charged leptons from the HNL decay that form a two-track displaced vertex. Even though the neutrino from the HNL decay is invisible to the detector, we can reconstruct the HNL mass using momentum conservation from the three reconstructed leptons and the information we have about the W mass distribution. The reconstructed HNL mass is then used to discriminate between signal and background.

Because most of the heavy flavour standard model background is short-lived, we can use a set of selection criteria to reduce the background from standard model particles to a negligible level. The remaining background comprises displaced vertices that are falsely reconstructed from two tracks that randomly cross in the tracker. A data-driven technique is used to estimate the random track crossing background.

For the first time, we set exclusion limits in the HNL coupling versus mass plane for both single-flavour and multi-flavour mixing scenarios motivated by neutrino-flavour oscillation results for both the normal and inverted neutrino-mass hierarchies. The model-dependent shape of the exclusion limits shows that limit summary plots have to be read cautiously and that more experiments should report their results in real physics multi-flavour mixing benchmark points that agree with the observed neutrino oscillation data. The benchmark points used in this work have been agreed upon in the feebly-interacting particles workshop 2020 at CERN [190]. The strongest limits are observed for a quasi-Dirac pair of HNLs with highly degenerate mass

and an equal mixing to all neutrino flavours of $1/3$ of the total mixing that agrees with recent neutrino oscillation data in the inverted neutrino mass hierarchy.

The ongoing search for long-lived HNLs that aims to analyse the full LHC Run2 proton-proton collision dataset (2015-2018, 13 TeV, 139 fb^{-1}) in the new software release and the partial LHC Run3 proton-proton collision dataset recorded between 2022 and 2023 with a centre-of-mass energy of 13.6 TeV and an integrated luminosity of 66 fb^{-1} is outlined. We benefit from an improved large-radius tracking algorithm to reconstruct the displaced HNL signatures and a more extensive set of available background simulations. Using background simulations with an enhanced data-driven background estimate allows us to tackle the low HNL mass region below 5 GeV, primarily excluded in the published analysis. In addition to the leptonic HNL decay, we continue our efforts to include the two-body semileptonic HNL decay into a charged lepton and a charged pion. Due to the branching ratio, the two-body semileptonic HNL decay would benefit the sensitivity in the low HNL mass region below 5 GeV. A second benefit of the two-body semileptonic HNL decay is the sharp mass peak in the invariant mass of the reconstructed displaced vertex, as both the charged lepton and the meson can be reconstructed. Additionally, we aim to increase the sensitivity to long-lived HNLs that decay outside the inner tracker by exploring a muon spectrometer-only vertexing technique. No results are shown because the search is ongoing at the time of writing. We aim to have a first publication draft by Summer 2024.

The search campaign for long-lived HNLs is one of the few pioneering dedicated long-lived particle searches in ATLAS. Together with the growing interest in those exotic signatures, the reconstruction algorithms and methods will improve. Long-lived particle searches play an important role in probing physics beyond the standard model.

8 OUTLOOK

The LHC Run3 is ongoing and will be followed by the upgrade of the LHC to the High Luminosity LHC (HL-LHC) in 2026. This upgrade will involve substantial detector upgrades for both ATLAS and CMS. The HL-LHC runs are expected to record an integrated luminosity of 3000-4000 fb⁻¹ of proton-proton collision data [98]. This extensive dataset will increase the sensitivity to HNL searches together with constantly evolving reconstruction algorithms and analysis techniques. The inclusion of long-lived particle-specific tracking algorithms in the standard reconstruction chain in ATLAS paves the way for more flexible analyses of those exotic signatures. Moving away from the standard view that interesting physics only occurs at the main interaction point opens up many opportunities. The interest in so-called feebly interacting particles is growing, and dedicated workshops that bring theorists and experimentalists together are happening regularly at CERN [190, 223]. It is crucial for those workshops to take place and for the community to agree on standard benchmark models to produce comparable results.

A near-future challenge would be to work on the implementation of the tau lepton channel, which has yet to be touched on in ATLAS. Including a tau lepton channel such as $W \rightarrow \tau \mathcal{N}$ and $\mathcal{N} \rightarrow \tau \tau \nu_\tau$ would not only enable us to probe the tau sector of the HNL coupling $|U_\tau|^2$, it would also boost the sensitivities for the Quasi-Dirac HNL model interpretation in both normal ($x_\tau = |U_\tau|^2/|U_{\text{tot}}|^2 = 0.48$) and inverted ($x_\tau = 1/3$) neutrino mass hierarchies. To achieve this, the reconstruction of displaced tau leptons would need to be implemented in ATLAS at low transverse momenta $p_T \lesssim 10$ GeV. Another challenge would be to trigger those events, as the single tau lepton triggers in ATLAS usually start at $p_T \gtrsim \mathcal{O}(100)$ GeV. A baseline could be to continue with single electron and muon triggers, which would only cover the leptonic tau lepton decay mode, or to use other triggers that combine signatures involving tau leptons.

The HL-LHC runs until 2040, but what comes next? The 2020 update of the European strategy for particle physics clearly states that an e^+e^- Higgs factory is the highest priority for the next collider [244], the US community study on the future of particle physics is on the same page [245]. One of those machines could be the Future Circular Collider (FCC)-ee [246], a proposed circular electron-positron collider with a circumference of about 100 km (90.7 km according to the overall lowest-risk baseline scenario [247]) located next to the LHC in Geneva as shown in fig. 8.1. A final feasibility study report is expected by the end of 2025. The FCC-ee would operate between 2045 and 2063 and could be upgraded to a $\sqrt{s} = 100$ TeV proton-proton collider (FCC-hh) with an operating time between 2070-2095. The next update of the European strategy for particle physics is foreseen for 2027 [222]. Competing options would be the proposed Circular




Figure 8.1: A schematic map showing a possible location for the Future Circular Collider (FCC) next to the Large Hadron Collider (LHC) at CERN. Image and modified caption are taken from ref. [222].




Figure 8.2: Constraints from past (shaded grey areas) and proposed future experiments (coloured lines) on the HNL coupling to muons only $|U_\mu|^2$ as a function of HNL mass m_N . The brown type I Seesaw mechanism bound gives a lower limit on the total coupling $|U|^2$ as a function of HNL mass to explain the observed pattern of neutrino flavour oscillations [41, 147]. The width of the brown type I Seesaw band indicates the range obtained by varying the lightest neutrino mass and light neutrino mass ordering. In contrast, the coloured bands give constraints on the total coupling $|U|^2$ as a function of HNL mass, surrounding the area that is consistent with the Seesaw mechanism and leptogenesis for two HNLs [63, 88] (inside the blue lines) and three HNLs [65] (inside the dark yellow lines) for the normal neutrino mass hierarchy (for vanishing (solid) and thermal (dotted) initial HNL abundances). The sensitivity estimates are given for the main detectors at the HL-LHC (red dotted lines) [196, 224–226] and proposed future colliders such as FCC [196, 227–229] (FCC-hh- η for detectors in $1.5 < |\eta| < 2.6$ that involve B decays [230]), LHeC and FCC-he [231], CEPC [232], ILC [233, 234], and a muon collider [235]. DV refers to dedicated displaced vertex searches. In the low-mass region, sensitivity estimates from SHiP [236, 237], MATHUSLA [238], and DUNE [239] dominate. Additional sensitivity estimates are shown for NA62 [240], FASER2 [241] and CODExb [242]. The light grey area on the lower left shows the Big Bang nucleosynthesis (BBN) constraint [149, 243] related to the abundance of matter over antimatter. Figure and modified caption taken from ref. [6].

Electron-Positron Collider (CEPC) [248] hosted by China, the International Linear Collider (ILC) [249] hosted by Japan, the Compact Linear Collider (CLIC) [250] at CERN, the Cool Copper Collider (C3) [251] and the Hybrid Asymmetric Linear Higgs Factory (HALHF) [252].

Another part of the European roadmap for accelerator research and development is a design study for a muon collider [253]. A muon collider operating at multi-TeV energies could serve as a discovery and precision instrument and would be an additional competing technology to a machine like the FCC. But it is essential to develop new technologies to account for the short lifetimes of muons.

Returning to HNLs, most of the new collider technologies would significantly push the sensitivity limits, as fig. 8.2 summarises [6]. Shown are several expected sensitivity limits in the HNL coupling versus mass plane for proposed future experiments, as well as constraints already set by past experiments. The type I Seesaw mechanism sets lower boundaries on the smallness of the HNL coupling values that agree with the lightest neutrino mass and the two possible neutrino mass hierarchies [41, 147]. Additional constraints surround the area in the coupling versus mass plane that is consistent with the type I Seesaw mechanism and leptogenesis for the case of two HNLs (blue band in fig. 8.2 [63, 88]) and three HNLs (dark yellow band in fig. 8.2 [65]). The bounds related to leptogenesis and the type I Seesaw mechanism for two HNLs have been detailed in fig. 3.6. A dedicated displaced vertex search of the HL-LHC (indicated as a red dotted line with the label 'LHC DV' in fig. 8.2) would enter the coupling versus mass region that agrees with the two HNL leptogenesis scenario at masses between $2 \lesssim m_{\mathcal{N}} \lesssim 32$ GeV. The HL-LHC will also push the higher mass and coupling area limits for prompt HNL searches. Future colliders, such as the FCC-ee or the CEPC, could reach masses up to 70 GeV at couplings close to the Seesaw limit (for $m_{\mathcal{N}} \approx 30$ GeV), while the proposed beam-dump experiment SHiP would be able to probe the low mass regime $0.2 \lesssim m_{\mathcal{N}} \lesssim 6$ GeV, which could be extended further to $0.2 \lesssim m_{\mathcal{N}} \lesssim 20$ GeV at low coupling values by experiments involving B decays at the FCC-hh (indicated as FCC-hh- η in fig. 8.2 [230]). The FCC-he, the ILC, the FCC-hh, and a muon collider could test the high mass region $0.1 \lesssim m_{\mathcal{N}} \lesssim 10$ TeV that would agree with leptogenesis scenarios involving three HNLs. But even with all the proposed experiments, the entire whitespace in fig. 8.2 cannot be explored. HNLs might be as heavy as $\sim 10^{14}$ GeV (see section 3.2); we don't know, but we can continue searching for HNLs within the limits of current and future experiments.

The future will be inspiring and full of new small-scale developments that could lead to discoveries. The open and curiosity-driven mentality at ATLAS and CERN will drive me personally to continue my search campaign for HNLs and physics beyond the standard model together with my colleagues at ATLAS.

The end.

A APPENDIX

A.1 LIKELIHOOD AND MAXIMUM LIKELIHOOD METHOD

A.1.1 PROBABILITY DENSITY FUNCTIONS

Consider a random variable x , which can be discrete or continuous. The probability prediction of x is given by a function $f(x)$. In the case of x being discrete, $f(x)$ itself is a probability. For the continuous case, $f(x)$ has the inverse dimension of x and $\int f(x)dx$ is the probability. $f(x)$ is the probability density function (PDF).

POISSON DISTRIBUTION

In a simple counting experiment, we count the number n of detected particles. They will be distributed around the mean value ν following a Poisson distribution.

$$\text{Pois}(n|\nu) = \frac{\nu^n}{n!} e^{-\nu} \quad (\text{A.1})$$

The standard deviation of the Poisson distribution follows with $\sigma = \sqrt{\nu}$.

GAUSSIAN DISTRIBUTION

$$\text{Gaus}(x|\mu, \sigma) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \quad (\text{A.2})$$

where μ is the mean or expectation value of the distribution, and σ is the standard deviation. To calculate the coverage probability of the σ bands in the limit plots, the integral of the Gaussian distribution is needed

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-t^2/2} dt \quad (\text{A.3})$$

The cumulative density function (CDF) for the Gaussian distribution reads

$$F(x) = \phi\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{2} \left[1 + \text{erf}\left(\frac{x-\mu}{\sqrt{2}\sigma}\right) \right] \quad (\text{A.4})$$

The probability p that a normal distributed quantity is inside the interval $[\mu - n\sigma, \mu + n\sigma]$ follows with

$$p = F(\mu + n\sigma) - F(\mu - n\sigma) = \operatorname{erf}\left[\frac{n}{\sqrt{2}}\right] \quad (\text{A.5})$$

The statistical significance α is calculated by

$$\alpha = 1 - p \quad (\text{A.6})$$

Important to mention is the central limit theorem, which states that every probability density distribution of finite variance convoluted many times with itself converges to a Gaussian probability density function [254].

A.1.2 LIKELIHOOD

The likelihood $\mathcal{L}(x|\theta)$ gives the probability of measuring x given some parameter θ . If multiple measurements x_i distribute according to different probability density functions $f_i(x_i|\theta)$, then the likelihood is given by

$$L(\mathbf{x}|\theta) = \prod_i f(x_i|\theta) \quad (\text{A.7})$$

The likelihood principle states that the likelihood function contains all information relevant for the measurement [254].

A.1.3 MAXIMUM LIKELIHOOD METHOD

Let us assume the measurements $\mathbf{x}$ depend on a set of unknown parameters θ . The values $\hat{\theta}$ for which the likelihood function $\mathcal{L}(\mathbf{x}|\theta)$ has a global maximum are called maximum likelihood estimate (MLE). To find the global maximum of the likelihood function, it is often convenient to use the negative logarithm of the likelihood function

$$-\ln \mathcal{L}(\mathbf{x}|\theta) = -\sum_i \ln f(x_i|\theta) \quad (\text{A.8})$$

A necessary condition for the minimum of the log-likelihood function follows with

$$-\frac{\partial \ln \mathcal{L}(\mathbf{x}|\hat{\theta})}{\partial \theta_j} = 0 \quad (\text{A.9})$$

where normalised probability density functions have to be used to build the likelihood such that

$$\int \mathcal{L}(\mathbf{x}|\theta) dx_1 dx_2 \cdots = 1 \quad (\text{A.10})$$

This minimisation is usually done numerically by frameworks like MINUIT [255]. After the maximum likelihood function has been found, the values for the previously unknown parameters $\hat{\theta}$ can be assigned. Since

the parameters $\hat{\theta}$ are found at the global maximum of the likelihood function, they represent the most likely set of parameters [254].

A.2 HNL MASS CALCULATION

This section is mainly taken from our PRL paper [10].

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

$$p_{\mathcal{N}}^\mu = p_2^\mu + p_3^\mu + p_\nu^\mu \equiv p_{23}^\mu + p_\nu^\mu. \quad (\text{A.11})$$

Four-momentum conservation in the W decay gives

$$p_W^\mu = p_1^\mu + p_{\mathcal{N}}^\mu = p_1^\mu + p_{23}^\mu + p_\nu^\mu. \quad (\text{A.12})$$

and we define

$$\begin{aligned} p_{23}^2 &= E_{23}^2 - |\vec{p}_{23}|^2 \equiv m_{23}^2 \\ p_{23}^\parallel &\equiv \vec{p}_{23} \cdot \hat{\mathbf{v}} \\ p_{23}^\perp &\equiv |\vec{p}_{23} - p_{23}^\parallel \hat{\mathbf{v}}| \end{aligned}$$

where m , E , and $|\vec{p}|$ are the mass, energy, and momentum-vector magnitude of the particles indicated by their subscript and $\hat{\mathbf{v}}$ 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 $k = (\hat{\mathbf{x}}', \hat{\mathbf{y}}', \hat{\mathbf{z}}')$, 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

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

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 un-

known variable α , which is the component of neutrino momentum in the $\hat{\mathbf{z}}'$ direction. To express eqs. (A.11) and (A.12) in terms of α , the following quantities are defined

$$\vec{p}'_{23} \equiv \vec{q} \quad (\text{A.13})$$

$$\vec{q} = (0, |\vec{p}_{23} \times \hat{\mathbf{z}}'| \equiv q_{\perp}, \vec{p}_{23} \cdot \hat{\mathbf{z}}' \equiv q_z) \quad (\text{A.14})$$

$$\vec{p}'_{\nu} = (0, -q_{\perp}, \alpha) \quad (\text{A.15})$$

$$E'_{\nu} = \sqrt{q_{\perp}^2 + \alpha^2}. \quad (\text{A.16})$$

Squaring eq. (A.12) gives

$$p_W'^2 = m_W^2 = m_1^2 + m_{\nu}^2 + m_{23}^2 + 2p'_1 \cdot (p'_{23} + p'_{\nu}) + 2p'_{23} \cdot p'_{\nu} \quad (\text{A.17})$$

where

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

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. (A.17) to solve for E_{ν}' gives

$$E'_{\nu} = A + B\alpha \quad (\text{A.18})$$

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. (A.18) from eq. (A.16) 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

$$\alpha = \frac{-AB \pm \sqrt{(B^2 - 1)q_{\perp}^2 + A^2}}{(B^2 - 1)}. \quad (\text{A.19})$$

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 the forward emission of the neutrino concerning 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. (A.19) depends on m_W . ATLAS has measured the W -boson pole mass to be $M_W = 80.370 \pm 0.019$ GeV [256]. This measurement is combined in Ref. [257] 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, if $m_W = M_W$ in eq. (A.19), there may be 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 ($m_{W,\text{med}}$). 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 ($m_{W,\text{min}}$) is computed. From eq. (A.17), the mass of the W boson is given by

$$m_W^2 = m_{23}^2 + 2(E'_1 E'_{23} + E'_\nu(E'_1 + E'_{23}) - p'_{1,z} q_z + q_\perp^2 - \alpha(p'_{1,z} + q_z)) \quad (\text{A.20})$$

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

$$\frac{d(m_W^2/2)}{d\alpha} = (E'_1 + E'_{23}) \frac{dE'_\nu}{d\alpha} - (p'_{1,z} + q_z) = 0. \quad (\text{A.21})$$

Using

$$\frac{dE'_\nu}{d\alpha} = \frac{d\sqrt{q_\perp^2 + \alpha^2}}{d\alpha} = \frac{\alpha}{E'_\nu}$$

in eq. (A.21), the chosen value of α that gives the minimum m_W is

$$\alpha = \frac{q_\perp B}{\sqrt{1 - B^2}}. \quad (\text{A.22})$$

Substituting eq. (A.22) into eq. (A.20), the minimum W boson mass is

$$m_{W,\text{min}}^2 = m_{23}^2 + 2 \left(E'_1 E'_{23} + (E'_1 + E'_{23}) \sqrt{q_\perp^2 + \frac{q_\perp^2 B^2}{1 - B^2}} - p'_{1,z} q_z + q_\perp^2 - (p'_{1,z} + q_z) \frac{q_\perp B}{\sqrt{1 - B^2}} \right).$$

The cumulative probability for the W boson to have a mass greater than $m_{W,\text{min}}$ 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

$$F(m_W^2) = \frac{1}{\pi} \arctan\left(\frac{m_W^2 - M_W^2}{M_W \Gamma_W}\right) + \frac{1}{2}. \quad (\text{A.23})$$

The midpoint of the allowed kinematic region has a value of

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

Rearranging eq. (A.23) for m_W^2 gives

$$m_W^2 = M_W^2 + \Gamma_W M_W \tan\left(\pi \left[F - \frac{1}{2}\right]\right). \quad (\text{A.24})$$

Substituting $F = F_{\text{med}}$ in eq. (A.24) gives an expression for the median W mass in the kinematically allowed region

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

This value of $m_{W,\text{med}}$ is used in eq. (A.19) to solve for α . From eq. (A.11) and the definitions in eqs. (A.13) to (A.16), the expression for the HNL mass in terms of α is

$$\begin{aligned} m_{\text{HNL}}^2 &= m_{23}^2 + 2p'_\nu \cdot p'_{23} \\ &= m_{23}^2 + 2E'_{23} \sqrt{q_\perp^2 + \alpha^2} + 2q_\perp^2 - 2q_z \alpha. \end{aligned} \quad (\text{A.25})$$

Substituting the expression for α in eq. (A.19) into eq. (A.25) gives the solution for the HNL mass.




Figure A.1: Truth reweighting validation p_T distributions for the four leptons (l_α , l_β , l_γ , ν_γ) between MADGRAPH5_AMC@NLO and PYTHIA for a $\mu-\mu\mu$ ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) sample. The applied weights are w_{CC+NC} in the 1SFH(μ) model ($x_e = 0$, $x_\mu = 1$, $x_\tau = 0$) for the lepton number violating decays (LNV). See section 5.4.5 for details. Figures are taken from our internal note ref. [16].

A.3 TRUTH REWEIGHTING VALIDATIONS

This section summarises the additional truth reweighting validation plots to account for the missing spin correlations in PYTHIA as described in section 5.4.5. The additional p_T distributions of the four leptons for the $\mu-\mu\mu$ channel and lepton number violating decays (fig. A.1), for the $\mu-\mu e$ channel and lepton number violating decays (fig. A.2), and for the $\mu-\mu e$ channel and lepton number conserving decays (fig. A.3) show good agreement between the reweighted PYTHIA and MADGRAPH5_AMC@NLO reference distributions.




Figure A.2: Truth reweighting validation p_T distributions for the four leptons (l_α , l_β , l_γ , ν_γ) between MADGRAPH5_AMC@NLO and PYTHIA for a $\mu-\mu e$ ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) sample. The applied weights are w_{CC} in the 1SFH(μ) model ($x_e = 0$, $x_\mu = 1$, $x_\tau = 0$) for the lepton number violating decays (LNV). See section 5.4.5 for details. Figures are taken from our internal note ref. [16].




Figure A.3: Truth reweighting validation p_T distributions for the four leptons (l_α , l_β , l_γ , ν_γ) between MADGRAPH5_AMC@NLO and PYTHIA for a $\mu\text{--}\mu e$ ($m_{\mathcal{N}} = 10$ GeV, $c\tau_{\mathcal{N}} = 10$ mm) sample. The applied weights are w_{CC} in the 1SFH(μ) model ($x_e = 0$, $x_\mu = 1$, $x_\tau = 0$) for the lepton number conserving decays (LNC). See section 5.4.5 for details. Figures are taken from our internal note ref. [16].

A.4 SIGNAL SELECTION EFFICIENCIES

This section summarises the signal selection efficiencies shown in fig. A.4 as a function of HNL mass and proper decay length for the four analysis channels ($e-ee$, $e-e\mu$, $\mu-ee$, $e-\mu\mu$) that are not shown in the main body in section 5.6.4. The selection efficiencies for the channels $\mu-\mu\mu$ and $\mu-\mu e$ are shown in fig. 5.24.




Figure A.4: The event selection efficiency for each mass-lifetime point for the remaining four analysis channels. The percentage of the produced MC simulation events that pass all signal region selections is shown. An entry of 0 indicates no events were selected [10]. The selection efficiencies for the channels $\mu-\mu\mu$ and $\mu-\mu e$ are shown in fig. 5.24. Figures are taken from our PRL publication ref. [10].

A.5 BACKGROUND UNCERTAINTIES IN FINAL BINNING

This section summarises the impact on the nominal data-driven background estimate distributions of the systematic uncertainties that arise from the data-driven background estimation method in the binning of the final fit (section 5.10.2) for the signal and control regions. The systematic uncertainties and the method used to derive the data-driven background estimate are explained in sections 5.7 and 5.9.2. The displaced vertex systematic takes the differences between opposite-sign and same-sign displaced vertices into account as described in section 5.9.2 and is summarised in fig. A.5 for the signal region and in fig. A.6 for the control region. The prompt lepton systematic is explained in section 5.9.2 and takes the finite number of prompt leptons available for shuffling into account. The impact of the nominal distribution of the prompt lepton systematic is shown in fig. A.7 for the signal region and in fig. A.8 for the control region.




Figure A.5: Single-bin HNL mass histogram of the data-driven shuffled background estimate (section 5.7.2) in the final signal region binning (one bin, $m_{\text{HNL}} \in [0, 20]$ GeV). The black histogram shows the nominal data-driven background estimate. The red and blue histograms indicate the symmetrised displaced vertex systematic uncertainty for $\pm 1\sigma$. The deviation of the systematic uncertainty in % from the nominal background estimate is shown at the bottom of each plot. The systematic uncertainties in the background estimate are fully correlated between the signal and control region.




Figure A.6: Single-bin HNL mass histogram of the data-driven shuffled background estimate (section 5.7.2) in the final control region binning (one bin, $m_{\text{HNL}} \in [20, 50]$ GeV). The black histogram shows the nominal data-driven background estimate. The red and blue histograms indicate the symmetrised displaced vertex systematic uncertainty for $\pm 1\sigma$. The deviation of the systematic uncertainty in % from the nominal background estimate is shown at the bottom of each plot. The systematic uncertainties in the background estimate are fully correlated between the signal and control region.




Figure A.7: Single-bin HNL mass histogram of the data-driven shuffled background estimate (section 5.7) in the final control region binning (one bin, $m_{\text{HNL}} \in [20, 50]$ GeV). The black histogram shows the nominal data-driven background estimate. The red and blue histograms indicate the symmetrised prompt lepton systematic uncertainty for $\pm 1\sigma$. The deviation of the systematic uncertainty in % from the nominal background estimate is shown at the bottom of each plot. The systematic uncertainties on the background estimate are fully correlated between the signal and control region.




Figure A.8: Single-bin HNL mass histogram of the data-driven shuffled background estimate (section 5.7) in the final signal region binning (one bin, $m_{\text{HNL}} \in [0, 20]$ GeV). The black histogram shows the nominal data-driven background estimate. The red and blue histograms indicate the symmetrised prompt lepton systematic uncertainty for $\pm 1\sigma$. The deviation of the systematic uncertainty in % from the nominal background estimate is shown at the bottom of each plot. The systematic uncertainties on the background estimate are fully correlated between the signal and control region.

A.6 TECHNICAL DETAILS OF THE DHNLLIMITS FRAMEWORK

This section describes the technical details of the DHNLLimits framework briefly, the statistical interpretation of the results itself is discussed in section 5.10.

The DHNLLimits framework is a python-based wrapper for the statistical analysis framework TREXFITTER [210] that is well established in ATLAS. The previous displaced HNL analysis [159] used a different framework unsuited for this work's scope. Therefore, the DHNLLimits framework was created from scratch after reviewing several existing statistical analysis frameworks. The framework is shown schematically in fig. A.9.

The DHNLLimits framework handles the job-creation for TREXFITTER and the required splitting, submission and book-keeping related to the parallelised CERN batch computation.

A standard TREXFITTER job involves seven different steps:

1. Input histogram production "n"
2. Pre-fit plot production "d"
3. Workspace creation "w"
4. Fitting "f"
5. Post-fit plot production "p"
6. Ranking plot production "r"
7. Limit-setting "l"

A job runs 1000 pseudo experiments for the baseline and the alternative hypothesis, and each tested signal strength value μ for a particular HNL mass, mean decay length, and HNL model (see section 5.10.1 for details). To reach the desired 10,000 pseudo experiments per mass, mean decay length, and model, ten jobs with different random seeds are run on the batch system in parallel and combined afterwards. The final analysis script then combines all $(m_{\mathcal{N}}, c\tau_{\mathcal{N}})$ points for a particular model and performs the cubic spline interpolation to obtain the final exclusion limit contours.




Figure A.9: Schematic illustration of the DHNLLimits framework. The Micro Ntuples (see also the schematic illustration of the analysis framework in fig. 5.17) are taken as an input for both background and signal MC samples. A python-steering script collects the inputs and splits the calculations done by TREXFITTER in small portions run in parallel on multiple computing nodes on the CERN batch system (or other national computing farms). The splitting is done in the HNL model, $c\tau_{\mathcal{N}}$, $m_{\mathcal{N}}$, and the number of pseudo experiments created. Each (model, $c\tau_{\mathcal{N}}$, $m_{\mathcal{N}}$) job is split into ten sub-jobs and combined afterwards to reduce the runtime for the toy MC experiments. A separate analysis script does the interpolation for the final contours.

A.7 IMPACT OF SYSTEMATIC UNCERTAINTIES ON THE EXPECTED 95% CONFIDENCE LIMIT

Figure A.10 shows the expected 95% confidence limit (CL) contours of the eight different models for the case where the systematic uncertainties are included in the maximum likelihood fit (solid lines) and for the case where only statistical uncertainties are taken into account (dashed lines). The excellent agreement between the expected 95% CL contours with and without systematic uncertainties shows that the sensitivity is mainly statistically limited. The systematic uncertainties on the simulated signal samples (section 5.9.1 together with the systematic uncertainties on the data-driven shuffled background estimate (section 5.9.2) do not noticeably restrict the sensitivity of the analysis.




Figure A.10: This figure shows the expected 95% CL contours for the eight different models in Dirac-limit (left) and Majorana-limit (right) including systematic uncertainties (solid line, 10k toy MC experiments) and excluding systematic uncertainties (dashed line, 'stat-only', 1k toy MC experiments).

BIBLIOGRAPHY

1. G. Bertone, D. Hooper, and J. Silk. “Particle dark matter: evidence, candidates and constraints”. *Physics Reports* 405:5, 2005, pp. 279–390. ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2004.08.031>. URL: <https://www.sciencedirect.com/science/article/pii/S0370157304003515>.
2. P.J. E. Peebles. “Growth of the nonbaryonic dark matter theory”. *Nature Astronomy* 1:3, 2017, p. 0057. DOI: <https://doi.org/10.1038/s41550-017-0057>. URL: <https://doi.org/10.1038/s41550-017-0057>.
3. R. Davis, D. S. Harmer, and K. C. Hoffman. “Search for Neutrinos from the Sun”. *Phys. Rev. Lett.* 20, 21 1968, pp. 1205–1209. DOI: [10.1103/PhysRevLett.20.1205](https://doi.org/10.1103/PhysRevLett.20.1205). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.20.1205>.
4. Y. Fukuda et al. “Evidence for Oscillation of Atmospheric Neutrinos”. *Phys. Rev. Lett.* 81, 8 1998, pp. 1562–1567. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.1562>.
5. Q. R. Ahmad et al. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. *Phys. Rev. Lett.* 89, 1 2002, p. 011301. DOI: [10.1103/PhysRevLett.89.011301](https://doi.org/10.1103/PhysRevLett.89.011301). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.89.011301>.
6. A. M. Abdullahi, P. Barham Alzás, B. Batell, J. Beacham, A. Boyarsky, S. Carbajal, A. Chatterjee, J. Crespo-Anadón, F. F. Deppisch, A. De Roeck, M. Drewes, A. Martin Gago, R. Gonzalez Suarez, E. Goudzovski, A. Hatzikoutelis, J. Hernandez-Garcia, M. Hostert, M. Hufnagel, P. Ilten, A. Izmaylov, K. J. Kelly, J. Klarić, J. Kopp, S. Kulkarni, M. Lamoureux, G. Lanfranchi, J. López-Pavón, O. Mikulenko, M. Mooney, M. Nemevšek, M. Ovchinnikov, S. Pascoli, R. Plestid, M. Rashad Darwish, F. Leo Redi, O. Ruchayskiy, R. Ruiz, M. Shaposhnikov, L. Shchutska, I. M. Shoemaker, R. Shrock, A. Sousa, N. Van Remortel, V. Syvolap, V. Takhistov, J.-L. Tastet, I. Timiryasov, A. C. Vincent, and J. Yu. “The present and future status of heavy neutral leptons”. *Journal of Physics G: Nuclear and Particle Physics* 50:2, 2023, p. 020501. DOI: [10.1088/1361-6471/ac98f9](https://doi.org/10.1088/1361-6471/ac98f9). URL: <https://dx.doi.org/10.1088/1361-6471/ac98f9>.
7. T. Asaka and M. Shaposhnikov. “The ν MSM, dark matter and baryon asymmetry of the universe”. *Physics Letters B* 620:1, 2005, pp. 17–26. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2005.06.020>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269305008087>.

8. T. Asaka, S. Blanchet, and M. Shaposhnikov. “The ν MSM, dark matter and neutrino masses”. *Physics Letters B* 631:4, 2005, pp. 151–156. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2005.09.070>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269305013675>.
9. F. Bezrukov, M. Y. Kalmykov, B. A. Kniehl, and M. Shaposhnikov. “Higgs boson mass and new physics”. *Journal of High Energy Physics* 2012:10, 2012, p. 140. DOI: [10.1007/JHEP10\(2012\)140](https://doi.org/10.1007/JHEP10(2012)140). URL: [https://doi.org/10.1007/JHEP10\(2012\)140](https://doi.org/10.1007/JHEP10(2012)140).
10. 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”. *Phys. Rev. Lett.* 131, 6 2023, p. 061803. DOI: [10.1103/PhysRevLett.131.061803](https://doi.org/10.1103/PhysRevLett.131.061803). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.061803>.
11. *The ATLAS Collaboration*. Accessed: 18-10-2023. URL: <https://atlas.cern/Discover/Collaboration>.
12. C. Appelt. *Internal: ATLAS ITk Qualification Task: development of a web application to summarize the status of a given component*. Technical report. Geneva: CERN, 2021. URL: <https://cds.cern.ch/record/2788956>.
13. *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*. Technical report. Geneva: CERN, 2017. DOI: [10.17181/CERN.FOZZ.ZP3Q](https://cds.cern.ch/record/2285585). URL: <https://cds.cern.ch/record/2285585>.
14. *Technical Design Report for the ATLAS Inner Tracker Strip Detector*. Technical report. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2257755>.
15. 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*. HEPData (collection). <https://doi.org/10.17182/hepdata.115511>. 2022.
16. D. A. Trischuk, C. Appelt, D. Casper, M. Danninger, S. Galantzan, S. Grancagnolo, H. Lackner, A. Lister, M. S. Lutz, M. Michetti, R. Newhouse, S. Oda, A. Soffer, and Y. Kehat. *Internal: Search for heavy neutral leptons produced in W boson decays using displaced vertices with the ATLAS detector*. Technical report. Geneva: CERN, 2020. URL: <https://cds.cern.ch/record/2714864>.
17. 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>.
18. R. Newhouse. “Using displaced tracks to search for sterile neutrinos in the ATLAS detector”. Presented 28 Nov 2022. University of British Columbia, 2022. URL: <https://cds.cern.ch/record/2845721>.
19. D. Dominguez. “Standard Model. Le modèle standard”, 2015. General Photo. URL: <https://cds.cern.ch/record/2002395>.

20. D. Griffiths. *Introduction to elementary particles*. John Wiley & Sons, 2020.
21. S. L. Glashow. “Partial-symmetries of weak interactions”. *Nuclear Physics* 22:4, 1961, pp. 579–588. ISSN: 0029-5582. DOI: [https://doi.org/10.1016/0029-5582\(61\)90469-2](https://doi.org/10.1016/0029-5582(61)90469-2). URL: <https://www.sciencedirect.com/science/article/pii/0029558261904692>.
22. S. Weinberg. “A Model of Leptons”. *Phys. Rev. Lett.* 19, 21 1967, pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.19.1264>.
23. A. Salam and J. C. Ward. “Weak and electromagnetic interactions”. *Il Nuovo Cimento (1955-1965)* 11:4, 1959, pp. 568–577. DOI: 10.1007/BF02726525. URL: <https://doi.org/10.1007/BF02726525>.
24. M. Peskin. *An introduction to quantum field theory*. CRC press, 2018.
25. P. Higgs. “Broken symmetries, massless particles and gauge fields”. *Physics Letters* 12:2, 1964, pp. 132–133. ISSN: 0031-9163. DOI: [https://doi.org/10.1016/0031-9163\(64\)91136-9](https://doi.org/10.1016/0031-9163(64)91136-9). URL: <https://www.sciencedirect.com/science/article/pii/0031916364911369>.
26. F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. *Phys. Rev. Lett.* 13, 9 1964, pp. 321–323. DOI: 10.1103/PhysRevLett.13.321. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.321>.
27. M. Thomson. *Modern particle physics*. Cambridge University Press, 2013.
28. N. Cabibbo. “Unitary Symmetry and Leptonic Decays”. *Phys. Rev. Lett.* 10, 12 1963, pp. 531–533. DOI: 10.1103/PhysRevLett.10.531. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.10.531>.
29. M. Kobayashi and T. Maskawa. “CP-Violation in the Renormalizable Theory of Weak Interaction”. *Progress of Theoretical Physics* 49:2, 1973, pp. 652–657. ISSN: 0033-068X. DOI: 10.1143/PTP.49.652. eprint: <https://academic.oup.com/ptp/article-pdf/49/2/652/5257692/49-2-652.pdf>. URL: <https://doi.org/10.1143/PTP.49.652>.
30. L. Wolfenstein. “Parametrization of the Kobayashi-Maskawa Matrix”. *Phys. Rev. Lett.* 51, 21 1983, pp. 1945–1947. DOI: 10.1103/PhysRevLett.51.1945. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.51.1945>.
31. Z. Maki, M. Nakagawa, and S. Sakata. “Remarks on the Unified Model of Elementary Particles”. *Progress of Theoretical Physics* 28:5, 1962, pp. 870–880. DOI: 10.1143/PTP.28.870. URL: <https://doi.org/10.1143/PTP.28.870>.
32. B. Pontecorvo. “Inverse beta processes and nonconservation of lepton charge”. *Zh. Eksp. Teor. Fiz.* 34, 1957, p. 247.
33. E. Noether. “Invariante Variationsprobleme”. ger. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1918, 1918, pp. 235–257. URL: <http://eudml.org/doc/59024>.

34. T. Nakano and K. Nishijima. “Charge Independence for V-particles”. *Progress of Theoretical Physics* 10:5, 1953, pp. 581–582. ISSN: 0033-068X. DOI: 10.1143/PTP.10.581. eprint: <https://academic.oup.com/ptp/article-pdf/10/5/581/5364926/10-5-581.pdf>. URL: <https://doi.org/10.1143/PTP.10.581>.
35. K. Nishijima. “Charge Independence Theory of V Particles”. *Progress of Theoretical Physics* 13:3, 1955, pp. 285–304. ISSN: 0033-068X. DOI: 10.1143/PTP.13.285. eprint: <https://academic.oup.com/ptp/article-pdf/13/3/285/5425869/13-3-285.pdf>. URL: <https://doi.org/10.1143/PTP.13.285>.
36. M. Gell-Mann. “The interpretation of the new particles as displaced charge multiplets”. *Il Nuovo Cimento (1955-1965)* 4:2, 1956, pp. 848–866. DOI: 10.1007/BF02748000. URL: <https://doi.org/10.1007/BF02748000>.
37. N. Arkani-Hamed, S. Dimopoulos, and G. Dvali. “The hierarchy problem and new dimensions at a millimeter”. *Physics Letters B* 429:3, 1998, pp. 263–272. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/S0370-2693\(98\)00466-3](https://doi.org/10.1016/S0370-2693(98)00466-3). URL: <https://www.sciencedirect.com/science/article/pii/S0370269398004663>.
38. *Standard Model Summary Plots February 2022*. Technical report. Geneva: CERN, 2022. URL: <https://cds.cern.ch/record/2804061>.
39. S. D. Bass, A. De Roeck, and M. Kado. “The Higgs boson implications and prospects for future discoveries”. *Nature Reviews Physics* 3:9, 2021, pp. 608–624. DOI: 10.1038/s42254-021-00341-2. URL: <https://doi.org/10.1038/s42254-021-00341-2>.
40. J. L. Tastet, O. Ruchayskiy, and I. Timiryasov. “Reinterpreting the ATLAS bounds on heavy neutral leptons in a realistic neutrino oscillation model”. *Journal of High Energy Physics* 2021:12, 2021, p. 182. DOI: 10.1007/JHEP12(2021)182. URL: [https://doi.org/10.1007/JHEP12\(2021\)182](https://doi.org/10.1007/JHEP12(2021)182).
41. I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou. “The fate of hints: updated global analysis of three-flavor neutrino oscillations”. *Journal of High Energy Physics* 2020:9, 2020, p. 178. DOI: 10.1007/JHEP09(2020)178. URL: [https://doi.org/10.1007/JHEP09\(2020\)178](https://doi.org/10.1007/JHEP09(2020)178).
42. KATRIN Collaboration. “Direct neutrino-mass measurement with sub-electronvolt sensitivity”. *Nature Physics* 18:2, 2022, pp. 160–166. DOI: 10.1038/s41567-021-01463-1. URL: <https://doi.org/10.1038/s41567-021-01463-1>.
43. S. Alam, M. Aubert, S. Avila, C. Balland, J. E. Bautista, M. A. Bershady, D. Bizyaev, M. R. Blanton, A. S. Bolton, J. Bovy, J. Brinkmann, J. R. Brownstein, E. Burtin, S. Chabanier, M. J. Chapman, P. D. Choi, C.-H. Chuang, J. Comparat, M.-C. Cousinou, A. Cuceu, K. S. Dawson, S. de la Torre, A. de Mattia, V. d. S. Agathe, H. d. M. des Bourbonx, S. Escoffier, T. Etourneau, J. Farr, A. Font-Ribera, P. M. Frinchaboy, S. Fromenteau, H. Gil-Marín, J.-M. Le Goff, A. X. Gonzalez-Morales, V. Gonzalez-Perez, K. Grabowski, J. Guy, A. J. Hawken, J. Hou, H. Kong, J. Parker, M. Klaene, J.-P. Kneib, S. Lin, D. Long, B. W. Lyke, A. de la Macorra, P. Martini, K. Masters, F. G. Mohammad, J. Moon, E.-M.

- Mueller, A. Muñoz Gutiérrez, A. D. Myers, S. Nadathur, R. Neveux, J. A. Newman, P. Noterdaeme, A. Oravetz, D. Oravetz, N. Palanque-Delabrouille, K. Pan, R. Paviot, W. J. Percival, I. Pérez-Ràfols, P. Petitjean, M. M. Pieri, A. Prakash, A. Raichoor, C. Ravoux, M. Rezaie, J. Rich, A. J. Ross, G. Rossi, R. Ruggeri, V. Ruhlmann-Kleider, A. G. Sánchez, F. J. Sánchez, J. R. Sánchez-Gallego, C. Sayres, D. P. Schneider, H.-J. Seo, A. Shafieloo, A. c. v. Slosar, A. Smith, J. Stermer, A. Tamone, J. L. Tinker, R. Tojeiro, M. Vargas-Magaña, A. Variu, Y. Wang, B. A. Weaver, A.-M. Weijmans, C. Yèche, P. Zarrouk, C. Zhao, G.-B. Zhao, and Z. Zheng. “Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory”. *Phys. Rev. D* 103, 8 2021, p. 083533. DOI: 10.1103/PhysRevD.103.083533. URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.083533>.
44. Planck Collaboration et al. “Planck 2018 results - VI. Cosmological parameters”. *A&A* 641, 2020, A6. DOI: 10.1051/0004-6361/201833910. URL: <https://doi.org/10.1051/0004-6361/201833910>.
 45. P. Minkowski. “ $\mu \rightarrow e\gamma$ at a rate of one out of 109 muon decays?” *Physics Letters B* 67:4, 1977, pp. 421–428. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(77\)90435-X](https://doi.org/10.1016/0370-2693(77)90435-X). URL: <https://www.sciencedirect.com/science/article/pii/037026937790435X>.
 46. T. Yanagida. “Horizontal gauge symmetry and masses of neutrinos”. *Conf. Proc. C* 7902131, 1979. Ed. by O. Sawada and A. Sugamoto, pp. 95–99.
 47. S. L. Glashow. “The Future of Elementary Particle Physics”. *NATO Sci. Ser. B* 61, 1980, p. 687. DOI: 10.1007/978-1-4684-7197-7_15.
 48. M. Gell-Mann, P. Ramond, and R. Slansky. “Complex Spinors and Unified Theories”. *Conf. Proc. C* 790927, 1979, pp. 315–321. arXiv: 1306.4669 [hep-th].
 49. R. N. Mohapatra and G. Senjanovic. “Neutrino Mass and Spontaneous Parity Nonconservation”. *Phys. Rev. Lett.* 44, 14 1980, pp. 912–915. DOI: 10.1103/PhysRevLett.44.912. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.44.912>.
 50. R. N. Mohapatra and G. Senjanovic. “Neutrino masses and mixings in gauge models with spontaneous parity violation”. *Phys. Rev. D* 23, 1 1981, pp. 165–180. DOI: 10.1103/PhysRevD.23.165. URL: <https://link.aps.org/doi/10.1103/PhysRevD.23.165>.
 51. J. Schechter and J. W. F. Valle. “Neutrino masses in $SU(2) \otimes U(1)$ theories”. *Phys. Rev. D* 22, 9 1980, pp. 2227–2235. DOI: 10.1103/PhysRevD.22.2227. URL: <https://link.aps.org/doi/10.1103/PhysRevD.22.2227>.
 52. J. Schechter and J. W. F. Valle. “Neutrino decay and spontaneous violation of lepton number”. *Phys. Rev. D* 25, 3 1982, pp. 774–783. DOI: 10.1103/PhysRevD.25.774. URL: <https://link.aps.org/doi/10.1103/PhysRevD.25.774>.

53. S. Weinberg. “Baryon- and Lepton-Nonconserving Processes”. *Phys. Rev. Lett.* 43, 21 1979, pp. 1566–1570. DOI: 10.1103/PhysRevLett.43.1566. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.43.1566>.
54. M. Magg and C. Wetterich. “Neutrino mass problem and gauge hierarchy”. *Physics Letters B* 94:1, 1980, pp. 61–64. DOI: [https://doi.org/10.1016/0370-2693\(80\)90825-4](https://doi.org/10.1016/0370-2693(80)90825-4). URL: <https://www.sciencedirect.com/science/article/pii/0370269380908254>.
55. R. Foot, H. Lew, X. G. He, and G. C. Joshi. “See-saw neutrino masses induced by a triplet of leptons”. *Zeitschrift für Physik C Particles and Fields* 44:3, 1989, pp. 441–444. DOI: 10.1007/BF01415558. URL: <https://doi.org/10.1007/BF01415558>.
56. E. Ma. “Pathways to Naturally Small Neutrino Masses”. *Phys. Rev. Lett.* 81, 6 1998, pp. 1171–1174. DOI: 10.1103/PhysRevLett.81.1171. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.1171>.
57. F. Vissani. “Do experiments suggest a hierarchy problem?” *Phys. Rev. D* 57, 11 1998, pp. 7027–7030. DOI: 10.1103/PhysRevD.57.7027. URL: <https://link.aps.org/doi/10.1103/PhysRevD.57.7027>.
58. J. A. Casas, J. Espinosa, and I. Hidalgo. “Implications for new physics from fine-tuning arguments 1. Application to SUSY and seesaw cases”. *Journal of High Energy Physics* 2004:11, 2004, p. 057. DOI: 10.1088/1126-6708/2004/11/057. URL: <https://dx.doi.org/10.1088/1126-6708/2004/11/057>.
59. J. D. Clarke, R. Foot, and R. R. Volkas. “Electroweak naturalness in the three-flavor type I seesaw model and implications for leptogenesis”. *Phys. Rev. D* 91, 7 2015, p. 073009. DOI: 10.1103/PhysRevD.91.073009. URL: <https://link.aps.org/doi/10.1103/PhysRevD.91.073009>.
60. I. Brivio and M. Trott. “Radiatively Generating the Higgs Potential and Electroweak Scale via the Seesaw Mechanism”. *Phys. Rev. Lett.* 119, 14 2017, p. 141801. DOI: 10.1103/PhysRevLett.119.141801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.141801>.
61. E. K. Akhmedov, V. A. Rubakov, and A. Y. Smirnov. “Baryogenesis via Neutrino Oscillations”. *Phys. Rev. Lett.* 81, 7 1998, pp. 1359–1362. DOI: 10.1103/PhysRevLett.81.1359. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.1359>.
62. M. Drewes, B. Garbrecht, P. Hernández, M. Kekic, J. Lopez-Pavon, J. Racker, N. Rius, J. Salvado, and D. Teresi. “ARS leptogenesis”. *International Journal of Modern Physics A* 33:05n06, 2018, p. 1842002. DOI: 10.1142/S0217751X18420022. URL: <https://doi.org/10.1142/S0217751X18420022>.
63. A. Granelli, K. Moffat, and S. T. Petcov. “Flavoured resonant leptogenesis at sub-TeV scales”. *Nuclear Physics B* 973, 2021, p. 115597. DOI: <https://doi.org/10.1016/j.nuclphysb.2021.115597>. URL: <https://www.sciencedirect.com/science/article/pii/S0550321321002947>.

64. J. Klaric, M. Shaposhnikov, and I. Timiryasov. “Reconciling resonant leptogenesis and baryogenesis via neutrino oscillations”. *Phys. Rev. D* 104, 5 2021, p. 055010. DOI: 10.1103/PhysRevD.104.055010. URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.055010>.
65. M. Drewes, Y. Georis, and J. Klaric. “Mapping the Viable Parameter Space for Testable Leptogenesis”. *Phys. Rev. Lett.* 128, 5 2022, p. 051801. DOI: 10.1103/PhysRevLett.128.051801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.128.051801>.
66. R. N. Mohapatra. “Mechanism for understanding small neutrino mass in superstring theories”. *Phys. Rev. Lett.* 56, 6 1986, pp. 561–563. DOI: 10.1103/PhysRevLett.56.561. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.56.561>.
67. M. B. Gavela, T. Hambye, D. Hernández, and P. Hernández. “Minimal flavour seesaw models”. *Journal of High Energy Physics* 2009:09, 2009, p. 038. DOI: 10.1088/1126-6708/2009/09/038. URL: <https://dx.doi.org/10.1088/1126-6708/2009/09/038>.
68. E. Akhmedov, M. Lindner, E. Schnapka, and J. F. Valle. “Left-right symmetry breaking in NJL approach”. *Physics Letters B* 368:4, 1996, pp. 270–280. DOI: [https://doi.org/10.1016/0370-2693\(95\)01504-3](https://doi.org/10.1016/0370-2693(95)01504-3). URL: <https://www.sciencedirect.com/science/article/pii/0370269395015043>.
69. E. Akhmedov, M. Lindner, E. Schnapka, and J. W. F. Valle. “Dynamical left-right symmetry breaking”. *Phys. Rev. D* 53, 5 1996, pp. 2752–2780. DOI: 10.1103/PhysRevD.53.2752. URL: <https://link.aps.org/doi/10.1103/PhysRevD.53.2752>.
70. M. Shaposhnikov. “A possible symmetry of the ν MSM”. *Nuclear Physics B* 763:1, 2007, pp. 49–59. DOI: <https://doi.org/10.1016/j.nuclphysb.2006.11.003>. URL: <https://www.sciencedirect.com/science/article/pii/S0550321306008960>.
71. J. Kersten and A. Y. Smirnov. “Right-handed neutrinos at CERN LHC and the mechanism of neutrino mass generation”. *Phys. Rev. D* 76, 7 2007, p. 073005. DOI: 10.1103/PhysRevD.76.073005. URL: <https://link.aps.org/doi/10.1103/PhysRevD.76.073005>.
72. A. Abada, C. Biggio, F. Bonnet, M. B. Gavela, and T. Hambye. “Low energy effects of neutrino masses”. *Journal of High Energy Physics* 2007:12, 2007, p. 061. DOI: 10.1088/1126-6708/2007/12/061. URL: <https://dx.doi.org/10.1088/1126-6708/2007/12/061>.
73. M. Drewes, B. Garbrecht, D. Gueter, and J. Klarić. “Testing the low scale seesaw and leptogenesis”. *Journal of High Energy Physics* 2017:8, 2017, p. 18. DOI: 10.1007/JHEP08(2017)018. URL: [https://doi.org/10.1007/JHEP08\(2017\)018](https://doi.org/10.1007/JHEP08(2017)018).
74. S. Dodelson and L. M. Widrow. “Sterile neutrinos as dark matter”. *Phys. Rev. Lett.* 72, 1 1994, pp. 17–20. DOI: 10.1103/PhysRevLett.72.17. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.72.17>.

75. X. Shi and G. M. Fuller. “New Dark Matter Candidate: Nonthermal Sterile Neutrinos”. *Phys. Rev. Lett.* 82, 14 1999, pp. 2832–2835. DOI: 10.1103/PhysRevLett.82.2832. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.82.2832>.
76. K. Abazajian, G. M. Fuller, and M. Patel. “Sterile neutrino hot, warm, and cold dark matter”. *Phys. Rev. D* 64, 2 2001, p. 023501. DOI: 10.1103/PhysRevD.64.023501. URL: <https://link.aps.org/doi/10.1103/PhysRevD.64.023501>.
77. T. Asaka, M. Shaposhnikov, and M. Laine. “Lightest sterile neutrino abundance within the ν MSM”. *Journal of High Energy Physics* 2007:01, 2007, p. 091. DOI: 10.1088/1126-6708/2007/01/091. URL: <https://dx.doi.org/10.1088/1126-6708/2007/01/091>.
78. M. Laine and M. Shaposhnikov. “Sterile neutrino dark matter as a consequence of ν MSM-induced lepton asymmetry”. *Journal of Cosmology and Astroparticle Physics* 2008:06, 2008, p. 031. DOI: 10.1088/1475-7516/2008/06/031. URL: <https://dx.doi.org/10.1088/1475-7516/2008/06/031>.
79. M. Escudero, N. Rius, and V. Sanz. “Sterile neutrino portal to Dark Matter II: exact dark symmetry”. *The European Physical Journal C* 77:6, 2017, p. 397. DOI: 10.1140/epjc/s10052-017-4963-x. URL: <https://doi.org/10.1140/epjc/s10052-017-4963-x>.
80. R. N. Mohapatra and J. C. Pati. ““Natural” left-right symmetry”. *Phys. Rev. D* 11, 9 1975, pp. 2558–2561. DOI: 10.1103/PhysRevD.11.2558. URL: <https://link.aps.org/doi/10.1103/PhysRevD.11.2558>.
81. G. Senjanovic and R. N. Mohapatra. “Exact left-right symmetry and spontaneous violation of parity”. *Phys. Rev. D* 12, 5 1975, pp. 1502–1505. DOI: 10.1103/PhysRevD.12.1502. URL: <https://link.aps.org/doi/10.1103/PhysRevD.12.1502>.
82. J. C. Pati and A. Salam. “Lepton number as the fourth “color””. *Phys. Rev. D* 10, 1 1974, pp. 275–289. DOI: 10.1103/PhysRevD.10.275. URL: <https://link.aps.org/doi/10.1103/PhysRevD.10.275>.
83. T. Asaka, S. Eijima, and H. Ishida. “Mixing of active and sterile neutrinos”. *Journal of High Energy Physics* 2011:4, 2011, p. 11. DOI: 10.1007/JHEP04(2011)011. URL: [https://doi.org/10.1007/JHEP04\(2011\)011](https://doi.org/10.1007/JHEP04(2011)011).
84. S. Davidson, G. Isidori, and A. Strumia. “The smallest neutrino mass”. *Physics Letters B* 646:2, 2007, pp. 100–104. DOI: <https://doi.org/10.1016/j.physletb.2007.01.015>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269307000871>.
85. J. A. Casas and A. Ibarra. “Oscillating neutrinos and $\mu \rightarrow e, \gamma$ ”. *Nuclear Physics B* 618:1, 2001, pp. 171–204. DOI: [https://doi.org/10.1016/S0550-3213\(01\)00475-8](https://doi.org/10.1016/S0550-3213(01)00475-8). URL: <https://www.sciencedirect.com/science/article/pii/S0550321301004758>.

86. A. Abada, S. Davidson, A. Ibarra, F.-X. Josse-Michaux, M. Losada, and A. Riotto. “Flavour matters in leptogenesis”. *Journal of High Energy Physics* 2006:09, 2006, p. 010. DOI: 10.1088/1126-6708/2006/09/010. URL: <https://dx.doi.org/10.1088/1126-6708/2006/09/010>.
87. S. Eijima, M. Shaposhnikov, and I. Timiryasov. “Parameter space of baryogenesis in the ν MSM”. *Journal of High Energy Physics* 2019:7, 2019, p. 77. DOI: 10.1007/JHEP07(2019)077. URL: [https://doi.org/10.1007/JHEP07\(2019\)077](https://doi.org/10.1007/JHEP07(2019)077).
88. J. Klaric, M. Shaposhnikov, and I. Timiryasov. “Uniting Low-Scale Leptogenesis Mechanisms”. *Phys. Rev. Lett.* 127, 11 2021, p. 111802. DOI: 10.1103/PhysRevLett.127.111802. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.111802>.
89. C. Service graphique. “Overall view of the LHC”, 2014. General Photo. URL: <https://cds.cern.ch/record/1708849>.
90. L. Evans. “The Large Hadron Collider”. *New Journal of Physics* 9:9, 2007, p. 335. DOI: 10.1088/1367-2630/9/9/335. URL: <https://dx.doi.org/10.1088/1367-2630/9/9/335>.
91. L. Evans and P. Bryant. “LHC Machine”. *Journal of Instrumentation* 3:08, 2008, S08001. DOI: 10.1088/1748-0221/3/08/S08001. URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08001>.
92. ATLAS Collaboration. “The ATLAS Experiment at the CERN Large Hadron Collider”. *Journal of Instrumentation* 3:08, 2008, S08003. DOI: 10.1088/1748-0221/3/08/S08003. URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08003>.
93. CMS Collaboration. “The CMS experiment at the CERN LHC”. *Journal of Instrumentation* 3:08, 2008, S08004. DOI: 10.1088/1748-0221/3/08/S08004. URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08004>.
94. LHCb Collaboration. “The LHCb Detector at the LHC”. *Journal of Instrumentation* 3:08, 2008, S08005. DOI: 10.1088/1748-0221/3/08/S08005. URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08005>.
95. ALICE Collaboration. “The ALICE experiment at the CERN LHC”. *Journal of Instrumentation* 3:08, 2008, S08002. DOI: 10.1088/1748-0221/3/08/S08002. URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08002>.
96. A. Lopes and M. L. Perrey. “FAQ-LHC The guide”. 2022. URL: <https://cds.cern.ch/record/2809109>.
97. ATLAS Collaboration. *ATLAS muon spectrometer: Technical Design Report*. Technical design report. CERN, 1997. URL: <https://cds.cern.ch/record/331068>.
98. *High lumi LHC plan*. Accessed: 14-10-2023. URL: <https://hilumilhc.web.cern.ch/article/ls3-schedule-change>.

99. L Arnaudon, P Baudrenghien, M Baylac, G Bellodi, Y Body, J Borburgh, P Bourquin, J Broere, O Brunner, L Bruno, C Carli, F. Caspers, S. M. Cousineau, Y Cuvet, C De Almeida Martins, T Dobers, T Fowler, R Garoby, F Gerigk, B Goddard, K Hanke, M Hori, M Jones, K Kahle, W. Kalbreier, T Kroyer, D K  chler, A. M. Lombardi, L. A. L  pez-Hernandez, M Magistris, M Martini, S Maury, E Page, M Paoluzzi, M Pasini, U Raich, C Rossi, J. P. Royer, E Sargsyan, J Serrano, R Scrivens, M Silari, M Timmins, W Venturini-Delsolaro, M Vretenar, R Wegner, W Weterings, and T Zickler. *Linac4 Technical Design Report*. Technical report. revised version submitted on 2006-12-14 09:00:40. Geneva: CERN, 2006. URL: <https://cds.cern.ch/record/1004186>.
100. K. Reich. *The CERN proton synchrotron booster*. Technical report. 1969.
101. E. Regenstreif. *The CERN proton synchrotron*. Technical report. CERN, 1959.
102. “The Super Proton Synchrotron”, 2012. URL: <https://cds.cern.ch/record/1997188>.
103. “Restarting the LHC: Why 13 TeV?”, 2014. URL: <https://cds.cern.ch/record/1998739>.
104. ATLAS Collaboration. *The ATLAS Experiment at the CERN Large Hadron Collider: A Description of the Detector Configuration for Run 3*. 2023. arXiv: 2305.16623 [physics.ins-det].
105. ATLAS IBL collaboration. “Production and integration of the ATLAS Insertable B-Layer”. *Journal of Instrumentation* 13:05, 2018, T05008. DOI: 10 . 1088 / 1748 - 0221 / 13 / 05 / T05008. URL: <https://dx.doi.org/10.1088/1748-0221/13/05/T05008>.
106.   kesson, Torsten and Eerola, Paula and Hedberg, Vincent and Jarlskog, G  ran and Lundberg, Bj  rn and Mj  rnmark, Ulf and Smirnova, Oxana and Almedhed, Sverker. *ATLAS computing: Technical Design Report*. eng. Technical report ATLAS TDR-017; CERN-LHCC-2005-022. CERN, 2005. URL: <https://lup.lub.lu.se/search/files/5842745/941311>.
107. J. Pequena   and P. Schaffner. “How ATLAS detects particles: diagram of particle paths in the detector”. 2013. URL: <https://cds.cern.ch/record/1505342>.
108. *ATLAS software documentation, tracking*. Accessed: 16-10-2023. URL: <https://atlassoftwaredocs.web.cern.ch/trackingTutorial/idooverview/>.
109. ATLAS Collaboration. “Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2”. *The European Physical Journal C* 77:10, 2017, p. 673. DOI: 10.1140/epjc/s10052-017-5225-7. URL: <https://doi.org/10.1140/epjc/s10052-017-5225-7>.
110. T Cornelissen, M Elsing, I Gavrilenko, W Liebig, E Moyse, and A Salzburger. “The new ATLAS track reconstruction (NEWT)”. *Journal of Physics: Conference Series* 119:3, 2008, p. 032014. DOI: 10.1088/1742-6596/119/3/032014. URL: <https://dx.doi.org/10.1088/1742-6596/119/3/032014>.

111. A. Rosenfeld and J. L. Pfaltz. “Sequential Operations in Digital Picture Processing”. *J. ACM* 13:4, 1966, pp. 471–494. ISSN: 0004-5411. DOI: 10.1145/321356.321357. URL: <https://doi.org/10.1145/321356.321357>.
112. E. Belau, R. Klanner, G. Lutz, E. Neugebauer, H. J. Seebrunner, A. Wylie, T. Böhringer, L. Hubbeling, P. Weilhammer, J. Kemmer, U. Kötz, and M. Riebesell. “Charge collection in silicon strip detectors”. *Nuclear Instruments and Methods in Physics Research* 214:2, 1983, pp. 253–260. DOI: [https://doi.org/10.1016/0167-5087\(83\)90591-4](https://doi.org/10.1016/0167-5087(83)90591-4). URL: <https://www.sciencedirect.com/science/article/pii/S0167508783905914>.
113. R. Frühwirth. “Application of Kalman filtering to track and vertex fitting”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 262:2, 1987, pp. 444–450. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/0168-9002\(87\)90887-4](https://doi.org/10.1016/0168-9002(87)90887-4). URL: <https://www.sciencedirect.com/science/article/pii/0168900287908874>.
114. ATLAS Collaboration. “A neural network clustering algorithm for the ATLAS silicon pixel detector”. *Journal of Instrumentation* 9:09, 2014, P09009. DOI: 10.1088/1748-0221/9/09/P09009. URL: <https://dx.doi.org/10.1088/1748-0221/9/09/P09009>.
115. A. Strandlie and R. Frühwirth. “Adaptive multitrack fitting”. *Computer Physics Communications* 133:1, 2000, pp. 34–42. DOI: [https://doi.org/10.1016/S0010-4655\(00\)00166-1](https://doi.org/10.1016/S0010-4655(00)00166-1). URL: <https://www.sciencedirect.com/science/article/pii/S0010465500001661>.
116. ATLAS Collaboration. *ATLAS software documentation, pile-up*. Accessed: 16-10-2023. URL: https://atlassoftwaredocs.web.cern.ch/AnalysisSWTutorial/cpalg_prw/.
117. P. Grafström and W. Kozanecki. “Luminosity determination at proton colliders”. *Progress in Particle and Nuclear Physics* 81, 2015, pp. 97–148. DOI: <https://doi.org/10.1016/j.pnpnp.2014.11.002>. URL: <https://www.sciencedirect.com/science/article/pii/S0146641014000878>.
118. *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*. Technical report. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2677054>.
119. W. Herr and B. Muratori. “Concept of luminosity”, 2006. DOI: 10.5170/CERN-2006-002.361. URL: <https://cds.cern.ch/record/941318>.
120. ATLAS Collaboration. *Public ATLAS luminosity plots*. Accessed: 16-10-2023. URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>.
121. G. Piacquadio, K. Prokofiev, and A. Wildauer. “Primary vertex reconstruction in the ATLAS experiment at LHC”. *Journal of Physics: Conference Series* 119:3, 2008, p. 032033. DOI: 10.1088/1742-6596/119/3/032033. URL: <https://dx.doi.org/10.1088/1742-6596/119/3/032033>.

122. ATLAS Collaboration. “Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC”. *The European Physical Journal C* 77:5, 2017, p. 332. DOI: 10.1140/epjc/s10052-017-4887-5. URL: <https://doi.org/10.1140/epjc/s10052-017-4887-5>.
123. ATLAS Collaboration. “Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV”. *Eur. Phys. J. C* 81, 2021, p. 578. DOI: 10.1140/epjc/s10052-021-09233-2. arXiv: 2012.00578 [hep-ex].
124. J. Illingworth and J. Kittler. “A survey of the hough transform”. *Computer Vision, Graphics, and Image Processing* 44:1, 1988, pp. 87–116. DOI: [https://doi.org/10.1016/S0734-189X\(88\)80033-1](https://doi.org/10.1016/S0734-189X(88)80033-1). URL: <https://www.sciencedirect.com/science/article/pii/S0734189X88800331>.
125. ATLAS Collaboration. “Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV.” *The European Physical Journal C* 79:8, 2019, p. 639. DOI: 10.1140/epjc/s10052-019-7140-6. URL: <https://doi.org/10.1140/epjc/s10052-019-7140-6>.
126. T. G. Cornelissen, M. Elsing, I. Gavrilenko, J.-F. Laporte, W. Liebig, M. Limper, K. Nikolopoulos, A. Poppleton, and A. Salzburger. “The global χ^2 track fitter in ATLAS”. *Journal of Physics: Conference Series* 119:3, 2008, p. 032013. DOI: 10.1088/1742-6596/119/3/032013. URL: <https://dx.doi.org/10.1088/1742-6596/119/3/032013>.
127. *Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung*. Technical report. Geneva: CERN, 2012. URL: <https://cds.cern.ch/record/1449796>.
128. R. Frühwirth. “A Gaussian-mixture approximation of the Bethe–Heitler model of electron energy loss by bremsstrahlung”. *Computer Physics Communications* 154:2, 2003, pp. 131–142. ISSN: 0010-4655. DOI: [https://doi.org/10.1016/S0010-4655\(03\)00292-3](https://doi.org/10.1016/S0010-4655(03)00292-3). URL: <https://www.sciencedirect.com/science/article/pii/S0010465503002923>.
129. ATLAS Collaboration. “Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016”. *The European Physical Journal C* 79:3, 2019, p. 205. DOI: 10.1140/epjc/s10052-019-6650-6. URL: <https://doi.org/10.1140/epjc/s10052-019-6650-6>.
130. M. Aaboud et al. “Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton-proton collision data”. *Journal of Instrumentation* 14:03, 2019, P03017. DOI: 10.1088/1748-0221/14/03/P03017. URL: <https://dx.doi.org/10.1088/1748-0221/14/03/P03017>.
131. A. Collaboration. *Electron and photon efficiencies in LHC Run 2 with the ATLAS experiment*. 2023. arXiv: 2308.13362 [hep-ex].

132. ATLAS Collaboration. “Search for displaced leptons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”. *Phys. Rev. Lett.* 127, 2020, p. 051802. DOI: 10.1103/PhysRevLett.127.051802. arXiv: 2011.07812 [hep-ex].
133. ATLAS Collaboration. *Performance of the reconstruction of large impact parameter tracks in the ATLAS inner detector*. Technical report. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2275635>.
134. ATLAS Collaboration. *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*. Technical report. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2669425>.
135. V Kostyukhin. *VKalVrt - package for vertex reconstruction in ATLAS*. Technical report. revised version number 1 submitted on 2003-09-24 11:10:53. Geneva: CERN, 2003. URL: <https://cds.cern.ch/record/685551>.
136. S. Höche. “Introduction to Parton-Shower Event Generators”. In: *Journeys Through the Precision Frontier: Amplitudes for Colliders*. 0. WORLD SCIENTIFIC, 2015, pp. 235–295. ISBN: 978-981-4678-75-9. DOI: doi:10.1142/9789814678766{_}0005. URL: https://doi.org/10.1142/9789814678766_0005.
137. L. A. Harland-Lang, A. D. Martin, P. Motylinski, and R. S. Thorne. “Parton distributions in the LHC era: MMHT 2014 PDFs”. *The European Physical Journal C* 75:5, 2015, p. 204. DOI: 10.1140/epjc/s10052-015-3397-6. URL: <https://doi.org/10.1140/epjc/s10052-015-3397-6>.
138. F James. “Monte Carlo theory and practice”. *Reports on Progress in Physics* 43:9, 1980, p. 1145. DOI: 10.1088/0034-4885/43/9/002. URL: <https://dx.doi.org/10.1088/0034-4885/43/9/002>.
139. G. T. Bodwin. “Factorization of the Drell-Yan cross section in perturbation theory”. *Phys. Rev. D* 31, 10 1985, pp. 2616–2642. DOI: 10.1103/PhysRevD.31.2616. URL: <https://link.aps.org/doi/10.1103/PhysRevD.31.2616>.
140. A. Buckley, J. Butterworth, S. Gieseke, D. Grellscheid, S. Höche, H. Hoeth, F. Krauss, L. Lönnblad, E. Nurse, P. Richardson, S. Schumann, M. H. Seymour, T. Sjöstrand, P. Skands, and B. Webber. “General-purpose event generators for LHC physics”. *Physics Reports* 504:5, 2011, pp. 145–233. DOI: <https://doi.org/10.1016/j.physrep.2011.03.005>. URL: <https://www.sciencedirect.com/science/article/pii/S0370157311000846>.
141. B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand. “Parton fragmentation and string dynamics”. *Physics Reports* 97:2, 1983, pp. 31–145. DOI: [https://doi.org/10.1016/0370-1573\(83\)90080-7](https://doi.org/10.1016/0370-1573(83)90080-7). URL: <https://www.sciencedirect.com/science/article/pii/0370157383900807>.

142. D. Amati and G. Veneziano. “Preconfinement as a property of perturbative QCD”. *Physics Letters B* 83:1, 1979, pp. 87–92. DOI: [https://doi.org/10.1016/0370-2693\(79\)90896-7](https://doi.org/10.1016/0370-2693(79)90896-7). URL: <https://www.sciencedirect.com/science/article/pii/0370269379908967>.
143. G. Aad et al. “The ATLAS Simulation Infrastructure”. *The European Physical Journal C* 70:3, 2010, pp. 823–874. DOI: [10.1140/epjc/s10052-010-1429-9](https://doi.org/10.1140/epjc/s10052-010-1429-9). URL: <https://doi.org/10.1140/epjc/s10052-010-1429-9>.
144. T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands. “An introduction to PYTHIA 8.2”. *Computer Physics Communications* 191, 2015, pp. 159–177. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2015.01.024>. URL: <https://www.sciencedirect.com/science/article/pii/S0010465515000442>.
145. S. Agostinelli et al. “Geant4—a simulation toolkit”. *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](https://doi.org/10.1016/S0168-9002(03)01368-8). URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688>.
146. W. Lukas. “Fast Simulation for ATLAS: Atlfast-II and ISF”. *Journal of Physics: Conference Series* 396:2, 2012, p. 022031. DOI: [10.1088/1742-6596/396/2/022031](https://doi.org/10.1088/1742-6596/396/2/022031). URL: <https://dx.doi.org/10.1088/1742-6596/396/2/022031>.
147. L. Canetti and M. Shaposhnikov. “Baryon asymmetry of the Universe in the ν MSM”. *Journal of Cosmology and Astroparticle Physics* 2010:09, 2010, p. 001. DOI: [10.1088/1475-7516/2010/09/001](https://doi.org/10.1088/1475-7516/2010/09/001). URL: <https://dx.doi.org/10.1088/1475-7516/2010/09/001>.
148. J Beacham, C Burrage, D Curtin, A De Roeck, J Evans, J. L. Feng, C Gatto, S Gninenko, A Hartin, I Irastorza, J Jaeckel, K Jungmann, K Kirch, F Kling, S Knapen, M Lamont, G Lanfranchi, C Lazzeroni, A Lindner, F Martinez-Vidal, M Moulson, N Neri, M Papucci, I Pedraza, K Petridis, M Pospelov, A Rozanov, G Ruoso, P Schuster, Y Semertzidis, T Spadaro, C Vallée, and G Wilkinson. “Physics beyond colliders at CERN: beyond the Standard Model working group report”. *Journal of Physics G: Nuclear and Particle Physics* 47:1, 2020, p. 010501. DOI: [10.1088/1361-6471/ab4cd2](https://doi.org/10.1088/1361-6471/ab4cd2). URL: <https://dx.doi.org/10.1088/1361-6471/ab4cd2>.
149. N. Sabti, A. Magalich, and A. Filimonova. “An extended analysis of Heavy Neutral Leptons during Big Bang Nucleosynthesis”. *Journal of Cosmology and Astroparticle Physics* 2020:11, 2020, p. 056. DOI: [10.1088/1475-7516/2020/11/056](https://doi.org/10.1088/1475-7516/2020/11/056). URL: <https://dx.doi.org/10.1088/1475-7516/2020/11/056>.
150. E. Fernandez-Martinez, J. Hernandez-Garcia, and J. Lopez-Pavon. “Global constraints on heavy neutrino mixing”. *Journal of High Energy Physics* 2016:8, 2016, p. 33. DOI: [10.1007/JHEP08\(2016\)033](https://doi.org/10.1007/JHEP08(2016)033). URL: [https://doi.org/10.1007/JHEP08\(2016\)033](https://doi.org/10.1007/JHEP08(2016)033).

151. M. Blennow, E. Fernández-Martínez, J. Hernández-García, J. López-Pavón, X. Marcano, and D. Naredo-Tuero. *Bounds on lepton non-unitarity and heavy neutrino mixing*. 2023. arXiv: 2306.01040 [hep-ph].
152. A. Aguilar-Arevalo, M. Aoki, M. Blecher, D. I. Britton, D. vom Bruch, D. A. Bryman, S. Chen, J. Comfort, S. Cuen-Rochin, L. Doria, P. Gumplinger, A. Hussein, Y. Igarashi, S. Ito, S. Kettell, L. Kurchaninov, L. S. Littenberg, C. Malbrunot, R. E. Mischke, T. Numao, D. Protopopescu, A. Sher, T. Sullivan, and D. Vavilov. “Improved search for heavy neutrinos in the decay $\pi \rightarrow e\nu$ ”. *Phys. Rev. D* 97, 7 2018, p. 072012. DOI: 10.1103/PhysRevD.97.072012. URL: <https://link.aps.org/doi/10.1103/PhysRevD.97.072012>.
153. D. A. Bryman and R. Shrock. “Constraints on sterile neutrinos in the MeV to GeV mass range”. *Phys. Rev. D* 100, 7 2019, p. 073011. DOI: 10.1103/PhysRevD.100.073011. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.073011>.
154. E. Cortina Gil et al. “Search for heavy neutral lepton production in K^+ decays to positrons”. *Physics Letters B* 807, 2020, p. 135599. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2020.135599>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269320304032>.
155. K. Abe et al. “Search for heavy neutrinos with the T2K near detector ND280”. *Phys. Rev. D* 100, 5 2019, p. 052006. DOI: 10.1103/PhysRevD.100.052006. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.052006>.
156. J. Dorenbosch, J. Allaby, U. Amaldi, G. Barbiellini, C. Berger, F. Bergsma, A. Capone, W. Flegel, L. Lanceri, M. Metcalf, C. Nieuwenhuis, J. Panman, K. Winter, I. Abt, J. Aspiazú, F. Büsler, H. Dautmann, P. Gall, T. Hebbeker, F. Niebergall, P. Schütt, P. Stähelin, P. Gorbunov, E. Grigoriev, V. Kافتanov, V. Khovansky, A. Rosanov, A. Baroncelli, L. Barone, B. Borgia, C. Bosio, M. Diemoz, U. Dore, F. Ferroni, E. Longo, L. Luminari, P. Monacelli, F. De Notaristefani, P. Pistilli, R. Santacesaria, C. Santoni, L. Tortora, and V. Valente. “A search for decays of heavy neutrinos in the mass range 0.5–2.8 GeV”. *Physics Letters B* 166:4, 1986, pp. 473–478. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(86\)91601-1](https://doi.org/10.1016/0370-2693(86)91601-1). URL: <https://www.sciencedirect.com/science/article/pii/0370269386916011>.
157. R. Barouki, G. Marocco, and S. Sarkar. “Blast from the past II: Constraints on heavy neutral leptons from the BEBC WA66 beam dump experiment”. *SciPost Phys.* 13, 2022, p. 118. DOI: 10.21468/SciPostPhys.13.5.118. URL: <https://scipost.org/10.21468/SciPostPhys.13.5.118>.
158. D. Collaboration. “Search for neutral heavy leptons produced in Z decays”. *Zeitschrift für Physik C Particles and Fields* 74:1, 1997, pp. 57–71. DOI: 10.1007/s002880050370. URL: <https://doi.org/10.1007/s002880050370>.

159. ATLAS Collaboration. “Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector”. *Journal of High Energy Physics* 2019:10, 2019, p. 265. DOI: 10.1007/JHEP10(2019)265. URL: [https://doi.org/10.1007/JHEP10\(2019\)265](https://doi.org/10.1007/JHEP10(2019)265).
160. CMS Collaboration. “Search for Heavy Neutral Leptons in Events with Three Charged Leptons in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV”. *Phys. Rev. Lett.* 120, 22 2018, p. 221801. DOI: 10.1103/PhysRevLett.120.221801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.221801>.
161. CMS Collaboration. “Search for long-lived heavy neutral leptons with displaced vertices in proton-proton collisions at $\sqrt{s} = 13$ TeV”. *Journal of High Energy Physics* 2022:7, 2022, p. 81. DOI: 10.1007/JHEP07(2022)081. URL: [https://doi.org/10.1007/JHEP07\(2022\)081](https://doi.org/10.1007/JHEP07(2022)081).
162. M. Gronau, C. N. Leung, and J. L. Rosner. “Extending limits on neutral heavy leptons”. *Phys. Rev. D* 29, 11 1984, pp. 2539–2558. DOI: 10.1103/PhysRevD.29.2539. URL: <https://link.aps.org/doi/10.1103/PhysRevD.29.2539>.
163. E. Fernández-Martínez, M. González-López, J. Hernández-García, M. Hostert, and J. López-Pavón. “Effective portals to heavy neutral leptons”, 2023. arXiv: 2304.06772 [hep-ph].
164. G. Bellini, J. Benziger, D. Bick, G. Bonfini, D. Bravo, M. Buizza Avanzini, B. Caccianiga, L. Cadenati, F. Calaprice, P. Cavalcante, A. Chavarria, A. Chepurinov, D. D’Angelo, S. Davini, A. Derbin, I. Drachnev, A. Empl, A. Etenko, K. Fomenko, D. Franco, C. Galbiati, S. Gazzana, C. Ghiano, M. Giammarchi, M. Göger-Neff, A. Goretti, L. Grandi, C. Hagner, E. Hungerford, A. Ianni, A. Ianni, V. Kobychiev, D. Korablev, G. Korga, D. Kryn, M. Laubenstein, T. Lewke, E. Litvinovich, B. Loer, F. Lombardi, P. Lombardi, L. Ludhova, G. Lukyanchenko, I. Machulin, S. Manecki, W. Maneschg, G. Manuzio, Q. Meindl, E. Meroni, L. Miramonti, M. Misiaszek, P. Mosteiro, V. Muratova, L. Oberauer, M. Obolensky, F. Ortica, K. Otis, M. Pallavicini, L. Papp, L. Perasso, S. Perasso, A. Pocar, G. Ranucci, A. Razeto, A. Re, A. Romani, N. Rossi, R. Saldanha, C. Salvo, S. Schönert, H. Simgen, M. Skorokhvatov, O. Smirnov, A. Sotnikov, S. Sukhotin, Y. Suvorov, R. Tartaglia, G. Testera, D. Vignaud, R. B. Vogelaar, F. von Feilitzsch, J. Winter, M. Wojcik, A. Wright, M. Wurm, J. Xu, O. Zaimidoroga, S. Zavatarelli, and G. Zuzel. “New limits on heavy sterile neutrino mixing in ^8B decay obtained with the Borexino detector”. *Phys. Rev. D* 88, 7 2013, p. 072010. DOI: 10.1103/PhysRevD.88.072010. URL: <https://link.aps.org/doi/10.1103/PhysRevD.88.072010>.
165. D. I. Britton, S. Ahmad, D. A. Bryman, R. A. Burnham, E. T. H. Clifford, P. Kitching, Y. Kuno, J. A. Macdonald, T. Numao, A. Olin, J. M. Poutissou, and M. S. Dixit. “Improved search for massive neutrinos in $\pi^+ \rightarrow e^+ \nu$ decay”. *Phys. Rev. D* 46, 3 1992, R885–R887. DOI: 10.1103/PhysRevD.46.R885. URL: <https://link.aps.org/doi/10.1103/PhysRevD.46.R885>.

166. D. Liventsev et al. “Search for heavy neutrinos at Belle”. *Phys. Rev. D* 87, 7 2013, p. 071102. DOI: 10.1103/PhysRevD.87.071102. URL: <https://link.aps.org/doi/10.1103/PhysRevD.87.071102>.
167. P. Achard et al. “Search for heavy isosinglet neutrino in e+e- annihilation at LEP”. *Physics Letters B* 517:1, 2001, pp. 67–74. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/S0370-2693\(01\)00993-5](https://doi.org/10.1016/S0370-2693(01)00993-5). URL: <https://www.sciencedirect.com/science/article/pii/S0370269301009935>.
168. M. Daum, P.-R. Kettle, B. Jost, R. M. Marshall, R. C. Minehart, W. A. Stephens, and K. O. H. Ziock. “Search for admixtures of massive neutrinos in the decay $\pi^+ \rightarrow \mu^+ + \nu$ ”. *Phys. Rev. D* 36, 9 1987, pp. 2624–2632. DOI: 10.1103/PhysRevD.36.2624. URL: <https://link.aps.org/doi/10.1103/PhysRevD.36.2624>.
169. A. Aguilar-Arevalo, M. Aoki, M. Blecher, D. Britton, D. vom Bruch, D. Bryman, S. Chen, J. Comfort, L. Doria, S. Cuen-Rochin, P. Gumplinger, A. Hussein, Y. Igarashi, S. Ito, S. Kettell, L. Kurchaninov, L. Littenberg, C. Malbrunot, R. Mischke, T. Numao, D. Protopopescu, A. Sher, T. Sullivan, and D. Vavilov. “Search for heavy neutrinos in $\pi \rightarrow \mu\nu$ decay”. *Physics Letters B* 798, 2019, p. 134980. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2019.134980>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269319307026>.
170. Y. Ema, Z. Liu, K.-F. Lyu, and M. Pospelov. “Heavy Neutral Leptons from Stopped Muons and Pions”. *Journal of High Energy Physics* 2023:8, 2023, p. 169. DOI: 10.1007/JHEP08(2023)169. URL: [https://doi.org/10.1007/JHEP08\(2023\)169](https://doi.org/10.1007/JHEP08(2023)169).
171. C. A. Argüelles, N. Foppiani, and M. Hostert. “Heavy neutral leptons below the kaon mass at hadroscopic neutrino detectors”. *Phys. Rev. D* 105, 9 2022, p. 095006. DOI: 10.1103/PhysRevD.105.095006. URL: <https://link.aps.org/doi/10.1103/PhysRevD.105.095006>.
172. K. J. Kelly and P. A. N. Machado. “MicroBooNE experiment, NuMI absorber, and heavy neutral leptons”. *Phys. Rev. D* 104, 5 2021, p. 055015. DOI: 10.1103/PhysRevD.104.055015. URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.055015>.
173. E. Cortina Gil et al. “Search for K+ decays to a muon and invisible particles”. *Physics Letters B* 816, 2021, p. 136259. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2021.136259>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269321001994>.
174. A. Cooper-Sarkar, S. Haywood, M. Parker, S. Sarkar, K. Barnham, P. Bostock, M. Faccini-Turluer, H. Grässler, J. Guy, P. Hulth, K. Hultqvist, U. Idschok, H. Klein, H. Kreutzmann, J. Krstic, M. Mobayyen, D. Morrison, B. Nellen, M. Talebzadeh, W. Venus, D. Vignaud, H. Wachsmuth, W. Wittek, and B. Wunsch. “Search for heavy neutrino decays in the BEBC beam dump experiment”. *Physics Letters B* 160:1, 1985, pp. 207–211. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(85\)91493-5](https://doi.org/10.1016/0370-2693(85)91493-5). URL: <https://www.sciencedirect.com/science/article/pii/0370269385914935>.

175. A. Vaitaitis, R. B. Drucker, J. Formaggio, S. Koutsoliotas, T. Adams, A. Alton, S. Avvakumov, L. de Barbaro, P. de Barbaro, R. H. Bernstein, A. Bodek, T. Bolton, J. Brau, D. Buchholz, H. Budd, L. Bugel, J. Conrad, B. T. Fleming, R. Frey, J. Goldman, M. Goncharov, D. A. Harris, R. A. Johnson, J. H. Kim, M. J. Lamm, W. Marsh, D. Mason, K. S. McFarland, C. McNulty, J. Monroe, D. Naples, P. Nienaber, A. Romosan, W. K. Sakumoto, H. Schellman, M. H. Shaevitz, P. Spentzouris, E. G. Stern, M. Vakili, V. Wu, U. K. Yang, J. Yu, G. P. Zeller, and E. D. Zimmerman. “Search for Neutral Heavy Leptons in a High-Energy Neutrino Beam”. *Phys. Rev. Lett.* 83, 24 1999, pp. 4943–4946. DOI: 10.1103/PhysRevLett.83.4943. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.83.4943>.
176. J. Badier, C. Bemporad, J. Boucrot, O. Callot, P. Charpentier, A. M. Cnops, M. Crozon, P. Delpierre, J. F. Detoeuf, G. R. Giannini, R. Hammarström, P. Lariccia, R. Lorenzi, J. Mas, A. Micheline, A. Tilquin, J. K. Walker, and N. Collaboration. “Mass and lifetime limits on new long-lived particles in 300 GeV/c π -interactions”. *Zeitschrift für Physik C Particles and Fields* 31:1, 1986, pp. 21–32. DOI: 10.1007/BF01559588. URL: <https://doi.org/10.1007/BF01559588>.
177. M. Dentler, Á. Hernández-Cabezudo, J. Kopp, P. Machado, M. Maltoni, I. Martinez-Soler, and T. Schwetz. “Updated global analysis of neutrino oscillations in the presence of eV-scale sterile neutrinos”. *Journal of High Energy Physics* 2018:8, 2018, p. 10. DOI: 10.1007/JHEP08(2018)010. URL: [https://doi.org/10.1007/JHEP08\(2018\)010](https://doi.org/10.1007/JHEP08(2018)010).
178. R. Plestid. “Luminous solar neutrinos. II. Mass-mixing portals”. *Phys. Rev. D* 104, 7 2021, p. 075028. DOI: 10.1103/PhysRevD.104.075028. URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.075028>.
179. I. Boiarska, A. Boyarsky, O. Mikulenko, and M. Ovchinnikov. “Constraints from the CHARM experiment on heavy neutral leptons with tau mixing”. *Phys. Rev. D* 104, 9 2021, p. 095019. DOI: 10.1103/PhysRevD.104.095019. URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.095019>.
180. J Orloff, A Rozanov, and C Santoni. “Limits on the mixing of tau neutrino to heavy neutrinos”. *Physics Letters B* 550:1, 2002, pp. 8–15. DOI: [https://doi.org/10.1016/S0370-2693\(02\)02769-7](https://doi.org/10.1016/S0370-2693(02)02769-7). URL: <https://www.sciencedirect.com/science/article/pii/S0370269302027697>.
181. R. Acciarri, C. Adams, J. Asaadi, B. Baller, V. Basque, F. Cavanna, A. de Gouvêa, R. S. Fitzpatrick, B. Fleming, P. Green, C. James, K. J. Kelly, I. Lepetic, X. Luo, O. Palamara, G. Scanavini, M. Soderberg, J. Spitz, A. M. Szelc, W. Wu, and T. Yang. “New Constraints on Tau-Coupled Heavy Neutral Leptons with Masses $m_N = 280 - 970$ MeV”. *Phys. Rev. Lett.* 127, 12 2021, p. 121801. DOI: 10.1103/PhysRevLett.127.121801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.121801>.
182. BaBar Collaboration. “Search for heavy neutral leptons using tau lepton decays at BaBar”. *Phys. Rev. D* 107, 5 2023, p. 052009. DOI: 10.1103/PhysRevD.107.052009. URL: <https://link.aps.org/doi/10.1103/PhysRevD.107.052009>.

183. M. Tanabashi et al. “Review of Particle Physics”. *Phys. Rev. D* 98, 3 2018, p. 030001. DOI: 10.1103/PhysRevD.98.030001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.030001>.
184. ATLAS Collaboration. “Identification and energy calibration of hadronically decaying tau leptons with the ATLAS experiment in pp collisions at $\sqrt{s} = 8$ TeV”. *The European Physical Journal C* 75:7, 2015, p. 303. DOI: 10.1140/epjc/s10052-015-3500-z. URL: <https://doi.org/10.1140/epjc/s10052-015-3500-z>.
185. ATLAS Collaboration. *Identification of hadronic tau lepton decays using neural networks in the ATLAS experiment*. Technical report. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2688062>.
186. ATLAS Collaboration. *Measurement of the tau lepton reconstruction and identification performance in the ATLAS experiment using pp collisions at $\sqrt{s} = 13$ TeV*. Technical report. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2261772>.
187. M. Dittmar, A. Santamaria, M. C. Gonzalez-Garcia, and J. W. F. Valle. “Production mechanisms and signatures of isosinglet neutral heavy leptons in Z^0 decays”. *Nuclear Physics B* 332:1, 1990, pp. 1–19. DOI: [https://doi.org/10.1016/0550-3213\(90\)90028-C](https://doi.org/10.1016/0550-3213(90)90028-C). URL: <https://www.sciencedirect.com/science/article/pii/055032139090028C>.
188. K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy. “Phenomenology of GeV-scale heavy neutral leptons”. *Journal of High Energy Physics* 2018:11, 2018, p. 32. DOI: 10.1007/JHEP11(2018)032. URL: [https://doi.org/10.1007/JHEP11\(2018\)032](https://doi.org/10.1007/JHEP11(2018)032).
189. ATLAS Collaboration. “Measurement of $W^\pm$ and Z -boson production cross sections in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *Physics Letters B* 759, 2016, pp. 601–621. DOI: <https://doi.org/10.1016/j.physletb.2016.06.023>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269316302763>.
190. P. Agrawal, M. Bauer, J. Beacham, A. Berlin, A. Boyarsky, S. Cebrian, X. Cid-Vidal, D. d’Enterria, A. De Roeck, M. Drewes, B. Echenard, M. Giannotti, G. F. Giudice, S. Gninenko, S. Gori, E. Goudzovski, J. Heeck, P. Hernandez, M. Hostert, I. G. Irastorza, A. Izmaylov, J. Jaeckel, F. Kahlhoefer, S. Knapen, G. Krnjaic, G. Lanfranchi, J. Monroe, V. I. M. Outchoorn, J. Lopez-Pavon, S. Pascoli, M. Pospelov, D. Redigolo, A. Ringwald, O. Ruchayskiy, J. Ruderman, H. Russell, J. Salfeld-Nebgen, P. Schuster, M. Shaposhnikov, L. Shchutska, J. Shelton, Y. Soreq, Y. Stadnik, J. Swallow, K. Tobioka, and Y. D. Tsai. “Feebly-interacting particles: FIPs 2020 workshop report”. *The European Physical Journal C* 81:11, 2021, p. 1015. DOI: 10.1140/epjc/s10052-021-09703-7. URL: <https://doi.org/10.1140/epjc/s10052-021-09703-7>.
191. ATLAS Pythia 8 tunes to 7 TeV data. Technical report. Geneva: CERN, 2014. URL: <https://cds.cern.ch/record/1966419>.

192. R. D. Ball, V. Bertone, S. Carrazza, C. S. Deans, L. Del Debbio, S. Forte, A. Guffanti, N. P. Hartland, J. I. Latorre, J. Rojo, and M. Ubiali. “Parton distributions with LHC data”. *Nuclear Physics B* 867:2, 2013, pp. 244–289. ISSN: 0550-3213. DOI: <https://doi.org/10.1016/j.nuclphysb.2012.10.003>. URL: <https://www.sciencedirect.com/science/article/pii/S0550321312005500>.
193. *The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie-Landschoff diffractive model*. Technical report. Geneva: CERN, 2016. URL: <https://cds.cern.ch/record/2206965>.
194. J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”. *Journal of High Energy Physics* 2014:7, 2014, p. 79. DOI: [10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079). URL: [https://doi.org/10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079).
195. C. Degrande, O. Mattelaer, R. Ruiz, and J. Turner. “Fully automated precision predictions for heavy neutrino production mechanisms at hadron colliders”. *Phys. Rev. D* 94, 5 2016, p. 053002. DOI: [10.1103/PhysRevD.94.053002](https://doi.org/10.1103/PhysRevD.94.053002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.94.053002>.
196. S. Pascoli, R. Ruiz, and C. Weiland. “Heavy neutrinos with dynamic jet vetoes: multilepton searches at $\sqrt{s} = 14, 27$, and 100 TeV”. *Journal of High Energy Physics* 2019:6, 2019, p. 49. DOI: [10.1007/JHEP06\(2019\)049](https://doi.org/10.1007/JHEP06(2019)049). URL: [https://doi.org/10.1007/JHEP06\(2019\)049](https://doi.org/10.1007/JHEP06(2019)049).
197. R. Ruiz. “Quantitative study on helicity inversion in Majorana neutrino decays at the LHC”. *Phys. Rev. D* 103, 1 2021, p. 015022. DOI: [10.1103/PhysRevD.103.015022](https://doi.org/10.1103/PhysRevD.103.015022). URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.015022>.
198. ATLAS Collaboration. “Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”. *Phys. Rev. D* 97, 5 2018, p. 052012. DOI: [10.1103/PhysRevD.97.052012](https://doi.org/10.1103/PhysRevD.97.052012). URL: <https://link.aps.org/doi/10.1103/PhysRevD.97.052012>.
199. L. Lee, C. Ohm, A. Soffer, and T.-T. Yu. “Collider searches for long-lived particles beyond the Standard Model”. *Progress in Particle and Nuclear Physics* 106, 2019, pp. 210–255. DOI: <https://doi.org/10.1016/j.pnpnp.2019.02.006>. URL: <https://www.sciencedirect.com/science/article/pii/S0146641019300109>.
200. F. J. Massey. “The Kolmogorov-Smirnov Test for Goodness of Fit”. *Journal of the American Statistical Association* 46:253, 1951, pp. 68–78. DOI: [10.1080/01621459.1951.10500769](https://doi.org/10.1080/01621459.1951.10500769). URL: <https://www.tandfonline.com/doi/abs/10.1080/01621459.1951.10500769>.
201. *Early Inner Detector Tracking Performance in the 2015 data at $\sqrt{s} = 13$ TeV*. Technical report. Geneva: CERN, 2015. URL: <https://cds.cern.ch/record/2110140>.

202. ATLAS Collaboration. “Search for exotic decays of the Higgs boson into long-lived particles in pp collisions at $\sqrt{s} = 13$ TeV using displaced vertices in the ATLAS inner detector”. *Journal of High Energy Physics* 2021:11, 2021, p. 229. DOI: 10.1007/JHEP11(2021)229. URL: [https://doi.org/10.1007/JHEP11\(2021\)229](https://doi.org/10.1007/JHEP11(2021)229).
203. ATLAS Collaboration. “Search for Displaced Leptons in $\sqrt{s} = 13$ TeV pp Collisions with the ATLAS Detector”. *Phys. Rev. Lett.* 127, 5 2021, p. 051802. DOI: 10.1103/PhysRevLett.127.051802. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.051802>.
204. G. Avoni, M. Bruschi, G. Cabras, D. Caforio, N. Dehghanian, A. Floderus, B. Giacobbe, F. Gianuzzi, F. Giorgi, P. Grafström, V. Hedberg, F. L. Manghi, S. Meneghini, J. Pinfold, E. Richards, C. Sbarra, N. S. Cesari, A. Sbrizzi, R. Soluk, G. Uccielli, S. Valentinetti, O. Viazlo, M. Villa, C. Vittori, R. Vuillermet, and A. Zoccoli. “The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS”. *Journal of Instrumentation* 13:07, 2018, P07017. DOI: 10.1088/1748-0221/13/07/P07017. URL: <https://dx.doi.org/10.1088/1748-0221/13/07/P07017>.
205. ATLAS Collaboration. “Performance of electron and photon triggers in ATLAS during LHC Run 2”. *The European Physical Journal C* 80:1, 2020, p. 47. DOI: 10.1140/epjc/s10052-019-7500-2. URL: <https://doi.org/10.1140/epjc/s10052-019-7500-2>.
206. ATLAS Collaboration. “Performance of the ATLAS muon triggers in Run 2”. *Journal of Instrumentation* 15:09, 2020, P09015. DOI: 10.1088/1748-0221/15/09/P09015. URL: <https://dx.doi.org/10.1088/1748-0221/15/09/P09015>.
207. ATLAS Collaboration. “Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data”. *Journal of Instrumentation* 14:12, 2019, P12006. DOI: 10.1088/1748-0221/14/12/P12006. URL: <https://dx.doi.org/10.1088/1748-0221/14/12/P12006>.
208. ATLAS Collaboration. “Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV”. *The European Physical Journal C* 76:5, 2016, p. 292. DOI: 10.1140/epjc/s10052-016-4120-y. URL: <https://doi.org/10.1140/epjc/s10052-016-4120-y>.
209. M. Davier, A. Höcker, and Z. Zhang. “The physics of hadronic tau decays”. *Rev. Mod. Phys.* 78, 4 2006, pp. 1043–1109. DOI: 10.1103/RevModPhys.78.1043. URL: <https://link.aps.org/doi/10.1103/RevModPhys.78.1043>.
210. ATLAS Collaboration. *TRExFitter*. Accessed: 26-09-2023. URL: <https://gitlab.cern.ch/TRExStats/TRExFitter/>.
211. A. L. Read. “Modified frequentist analysis of search results (the CL_s method)”, 2000. DOI: 10.5170/CERN-2000-005.81. URL: <https://cds.cern.ch/record/451614>.

212. A. L. Read. “Presentation of search results: the CLs technique”. *Journal of Physics G: Nuclear and Particle Physics* 28:10, 2002, p. 2693. DOI: 10.1088/0954-3899/28/10/313. URL: <https://dx.doi.org/10.1088/0954-3899/28/10/313>.
213. N. Jerzy and P. E. Sharpe. “On the problem of the most efficient tests of statistical hypotheses”. *Philosophical Transactions of the Royal Society of London* 231:694-706, 1933. DOI: <https://doi.org/10.1098%2Frsta.1933.0009>.
214. S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. *The Annals of Mathematical Statistics* 9:1, 1938, pp. 60–62. DOI: 10.1214/aoms/1177732360. URL: <https://doi.org/10.1214/aoms/1177732360>.
215. G. Cowan, K. Cranmer, E. Gross, and O. Vitells. “Asymptotic formulae for likelihood-based tests of new physics”. *The European Physical Journal C* 71:2, 2011, p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. URL: <https://doi.org/10.1140/epjc/s10052-011-1554-0>.
216. *Formulae for Estimating Significance*. Technical report. Geneva: CERN, 2020. URL: <https://cds.cern.ch/record/2736148>.
217. A. Collaboration. *Performance of the reconstruction of large impact parameter tracks in the inner detector of ATLAS*. 2023. arXiv: 2304.12867 [hep-ex].
218. CMS Collaboration. *Search for long-lived heavy neutral leptons with lepton flavour conserving or violating decays to a jet and an electron, muon, or tau lepton*. Technical report. Geneva: CERN, 2023. URL: <https://cds.cern.ch/record/2852843>.
219. CMS Collaboration. *Search for long-lived heavy neutral leptons decaying in the CMS muon detectors in proton-proton collisions at $\sqrt{s} = 13$ TeV*. Technical report. Geneva: CERN, 2023. URL: <https://cds.cern.ch/record/2865227>.
220. ATLAS Collaboration. “Search for events with a pair of displaced vertices from long-lived neutral particles decaying into hadronic jets in the ATLAS muon spectrometer in pp collisions at $\sqrt{s} = 13$ TeV”. *Phys. Rev. D* 106, 3 2022, p. 032005. DOI: 10.1103/PhysRevD.106.032005. URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.032005>.
221. ATLAS Collaboration. *Public ATLAS luminosity plots Run3*. Accessed: 24-10-2023. URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3>.
222. *future circular collider*. Accessed: 27-10-2023. URL: <https://home.cern/science/accelerators/future-circular-collider>.

223. C. Antel, M. Battaglieri, J. Beacham, C. Boehm, O. Buchmüller, F. Calore, P. Carenza, B. Chauhan, P. Cladè, P. Coloma, P. Crivelli, V. Dandoy, L. Darmé, B. Dey, F. F. Deppisch, A. D. Roeck, M. Drewes, B. Echenard, V. V. Flambaum, P. Foldenauer, C. Gatti, M. Giannotti, A. Golutvin, M. C. Gonzalez-Garcia, S. Gori, E. Goudzovski, A. Granelli, H. Grote, S. Guellati-Khelifa, J. Hajer, P. Harris, C. Hearty, D. Heuchel, M. Hostert, S. Junius, F. Kahlhoefer, J. Klaric, F. Kling, P. Klose, J. Knolle, J. Kopp, O. Kwon, O. Lantwin, G. Lanfranchi, L. Li, A. Lindner, J. Lopez-Pavon, J. Marocco, J. W. Martin, S. Middleton, S. Milstead, I. Oceano, C. A. J. O'Hare, A. Paoloni, S. Pascoli, S. T. Petcov, M. Pospelov, R. Pöttgen, M. Raggi, G. Ripellino, I. B. Samsonov, S. Sandner, S. Söldner-Rembold, J. Shelton, N. Song, C. Sun, Y. V. Stadnik, J. L. Tastet, N. Toro, N. Tran, N. Trevisani, S. Ulmer, S. Urrea, B. Velghe, B. Wallisch, Y. Y. Y. Wong, C. Zorbilmez, and K. Zurek. *Feebly Interacting Particles: FIPs 2022 workshop report*. 2023. arXiv: 2305.01715 [hep-ph].
224. M. Drewes and J. Hajer. “Heavy neutrinos in displaced vertex searches at the LHC and HL-LHC”. *Journal of High Energy Physics* 2020:2, 2020, p. 70. DOI: 10.1007/JHEP02(2020)070. URL: [https://doi.org/10.1007/JHEP02\(2020\)070](https://doi.org/10.1007/JHEP02(2020)070).
225. E. Izaguirre and B. Shuve. “Multilepton and lepton jet probes of sub-weak-scale right-handed neutrinos”. *Phys. Rev. D* 91, 9 2015, p. 093010. DOI: 10.1103/PhysRevD.91.093010. URL: <https://link.aps.org/doi/10.1103/PhysRevD.91.093010>.
226. A. Das, P. Konar, and A. Thalapillil. “Jet substructure shedding light on heavy Majorana neutrinos at the LHC”. *Journal of High Energy Physics* 2018:2, 2018, p. 83. DOI: 10.1007/JHEP02(2018)083. URL: [https://doi.org/10.1007/JHEP02\(2018\)083](https://doi.org/10.1007/JHEP02(2018)083).
227. C. B. Verhaaren, J. Alimena, M. Bauer, P. Azzi, R. Ruiz, M. Neubert, O. Mikulenko, M. Ovchinnikov, M. Drewes, J. Klaric, A. Blondel, C. Rizzi, A. Sfyrta, T. Sharma, S. Kulkarni, A. Thamm, A. Blondel, R. G. Suarez, and L. Rygaard. “Searches for long-lived particles at the future FCC-ee”. *Frontiers in Physics* 10, 2022. ISSN: 2296-424X. DOI: 10.3389/fphy.2022.967881. URL: <https://www.frontiersin.org/articles/10.3389/fphy.2022.967881>.
228. S. Antusch, E. Cazzato, and O. Fischer. “Sterile neutrino searches at future e-e+, pp and e-p colliders”. *International Journal of Modern Physics A* 32:14, 2017, p. 1750078. DOI: 10.1142/S0217751X17500786. eprint: <https://doi.org/10.1142/S0217751X17500786>. URL: <https://doi.org/10.1142/S0217751X17500786>.
229. S. Antusch, E. Cazzato, O. Fischer, A. Hammad, and K. Wang. “Lepton flavor violating dilepton dijet signatures from sterile neutrinos at proton colliders”. *Journal of High Energy Physics* 2018:10, 2018, p. 67. DOI: 10.1007/JHEP10(2018)067. URL: [https://doi.org/10.1007/JHEP10\(2018\)067](https://doi.org/10.1007/JHEP10(2018)067).

230. A. Boyarsky, O. Mikulenko, M. Ovchinnikov, and L. Shchutska. “Exploring the potential of FCC-hh to search for particles from B mesons”. *Journal of High Energy Physics* 2023:1, 2023, p. 42. DOI: 10.1007/JHEP01(2023)042. URL: [https://doi.org/10.1007/JHEP01\(2023\)042](https://doi.org/10.1007/JHEP01(2023)042).
231. S. Antusch, O. Fischer, and A. Hammad. “Lepton-trijet and displaced vertex searches for heavy neutrinos at future electron-proton colliders”. *Journal of High Energy Physics* 2020:3, 2020, p. 110. DOI: 10.1007/JHEP03(2020)110. URL: [https://doi.org/10.1007/JHEP03\(2020\)110](https://doi.org/10.1007/JHEP03(2020)110).
232. Y.-F. Shen, J.-N. Ding, and Q. Qin. “Monojet search for heavy neutrinos at future Z-factories”. *The European Physical Journal C* 82:5, 2022, p. 398. DOI: 10.1140/epjc/s10052-022-10301-4. URL: <https://doi.org/10.1140/epjc/s10052-022-10301-4>.
233. K. Mękała, J. Reuter, and A. F. Żarnecki. “Heavy neutrinos at future linear e+e-colliders”. *Journal of High Energy Physics* 2022:6, 2022, p. 10. DOI: 10.1007/JHEP06(2022)010. URL: [https://doi.org/10.1007/JHEP06\(2022\)010](https://doi.org/10.1007/JHEP06(2022)010).
234. S. Antusch, E. Cazzato, and O. Fischer. “Displaced vertex searches for sterile neutrinos at future lepton colliders”. *Journal of High Energy Physics* 2016:12, 2016, p. 7. DOI: 10.1007/JHEP12(2016)007. URL: [https://doi.org/10.1007/JHEP12\(2016\)007](https://doi.org/10.1007/JHEP12(2016)007).
235. K. M. Black et al. *Muon Collider Forum Report*. 2023. arXiv: 2209.01318 [hep-ex].
236. C. Ahdida et al. “Sensitivity of the SHiP experiment to Heavy Neutral Leptons”. *Journal of High Energy Physics* 2019:4, 2019, p. 77. DOI: 10.1007/JHEP04(2019)077. URL: [https://doi.org/10.1007/JHEP04\(2019\)077](https://doi.org/10.1007/JHEP04(2019)077).
237. D. Gorbunov, I. Krasnov, Y. Kudenko, and S. Suvorov. “Heavy neutral leptons from kaon decays in the SHiP experiment”. *Physics Letters B* 810, 2020, p. 135817. ISSN: 0370-2693. DOI: <https://doi.org/10.1016/j.physletb.2020.135817>. URL: <https://www.sciencedirect.com/science/article/pii/S0370269320306201>.
238. D. Curtin, M. Drewes, M. McCullough, P. Meade, R. N. Mohapatra, J. Shelton, B. Shuve, E. Accomando, C. Alpigiani, S. Antusch, J. C. Arteaga-Velázquez, B. Batell, M. Bauer, N. Blinov, K. S. Caballero-Mora, J. H. Chang, E. J. Chun, R. T. Co, T. Cohen, P. Cox, N. Craig, C. Csáki, Y. Cui, F. D’Eramo, L. D. Rose, P. S. B. Dev, K. R. Dienes, J. A. Dror, R. Essig, J. A. Evans, J. L. Evans, A. F. Tellez, O. Fischer, T. Flacke, A. Fradette, C. Frugiuele, E. Fuchs, T. Gherghetta, G. F. Giudice, D. Gorbunov, R. S. Gupta, C. Hagedorn, L. J. Hall, P. Harris, J. C. Helo, M. Hirsch, Y. Hochberg, A. Hook, A. Ibarra, S. Ipek, S. Jung, S. Knapen, E. Kuflik, Z. Liu, S. Lombardo, H. J. Lubatti, D. McKeen, E. Molinaro, S. Moretti, N. Nagata, M. Neubert, J. M. No, E. Olaiya, G. Perez, M. E. Peskin, D. Pinner, M. Pospelov, M. Reece, D. J. Robinson, M. R. Cahuantzi, R. Santonico, M. Schlaffer, C. H. Shepherd-Themistocleous, A. Spray, D. Stolarski, M. A. S. Vasquez, R. Sundrum, A. Thamm, B. Thomas, Y. Tsai, B. Tweedie, S. M. West, C. Young, F. Yu, B. Zaldivar, Y. Zhang, K. Zurek, and J. Zurita. “Long-

- lived particles at the energy frontier: the MATHUSLA physics case”. *Reports on Progress in Physics* 82:11, 2019, p. 116201. DOI: 10.1088/1361-6633/ab28d6. URL: <https://dx.doi.org/10.1088/1361-6633/ab28d6>.
239. P. Ballett, T. Boschi, and S. Pascoli. “Heavy Neutral Leptons from low-scale seesaws at the DUNE Near Detector”. *Journal of High Energy Physics* 2020:3, 2020, p. 111. DOI: 10.1007/JHEP03(2020)111. URL: [https://doi.org/10.1007/JHEP03\(2020\)111](https://doi.org/10.1007/JHEP03(2020)111).
 240. M. Drewes, J. Hajer, J. Klaric, and G. Lanfranchi. “NA62 sensitivity to heavy neutral leptons in the low scale seesaw model”. *Journal of High Energy Physics* 2018:7, 2018, p. 105. DOI: 10.1007/JHEP07(2018)105. URL: [https://doi.org/10.1007/JHEP07\(2018\)105](https://doi.org/10.1007/JHEP07(2018)105).
 241. A. Ariga, T. Ariga, J. Boyd, F. Cadoux, D. W. Casper, Y. Favre, J. L. Feng, D. Ferrere, I. Galon, S. Gonzalez-Sevilla, S.-C. Hsu, G. Iacobucci, E. Kajomovitz, F. Kling, S. Kuehn, L. Levinson, H. Otono, B. Petersen, O. Sato, M. Schott, A. Sfyrta, J. Smolinsky, A. M. Soffa, Y. Takubo, E. Torrence, S. Trojanowski, and G. Zhang. “FASER’s physics reach for long-lived particles”. *Phys. Rev. D* 99, 9 2019, p. 095011. DOI: 10.1103/PhysRevD.99.095011. URL: <https://link.aps.org/doi/10.1103/PhysRevD.99.095011>.
 242. G. Aielli, E. Ben-Haim, R. Cardarelli, M. J. Charles, X. Cid Vidal, V. Coco, B. Dey, R. Dumps, J. A. Evans, G. Gibbons, O. L. Dortz, V. V. Gligorov, P. Ilten, S. Knapen, J. Lee, S. López Soliño, B. Nachman, M. Papucci, F. Polci, R. Quessard, H. Ramani, D. J. Robinson, H. Schindler, M. D. Sokoloff, P. Swallow, R. Vari, N. Watson, and M. Williams. “Expression of interest for the CODEX-b detector”. *The European Physical Journal C* 80:12, 2020, p. 1177. DOI: 10.1140/epjc/s10052-020-08711-3. URL: <https://doi.org/10.1140/epjc/s10052-020-08711-3>.
 243. A. Boyarsky, M. Ovchinnikov, O. Ruchayskiy, and V. Syvolap. “Improved big bang nucleosynthesis constraints on heavy neutral leptons”. *Phys. Rev. D* 104, 2 2021, p. 023517. DOI: 10.1103/PhysRevD.104.023517. URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.023517>.
 244. *2020 Update of the European Strategy for Particle Physics (Brochure)*. Technical report. Geneva, 2020. DOI: 10.17181/CERN.JSC6.W89E. URL: <https://cds.cern.ch/record/2721370>.
 245. J. N. Butler, R. S. Chivukula, A. de Gouvêa, T. Han, Y.-K. Kim, P. Cushman, G. R. Farrar, Y. G. Kolomensky, S. Nagaitsev, N. Yunes, S. Gourlay, T. Raubenheimer, V. Shiltsev, K. A. Assamagan, B. Quinn, V. D. Elvira, S. Gottlieb, B. Nachman, A. S. Chou, M. Soares-Santos, T. M. P. Tait, M. Narain, L. Reina, A. Tricoli, P. S. Barbeau, P. Merkel, J. Zhang, P. Huber, K. Scholberg, E. Worcester, M. Artuso, R. H. Bernstein, A. A. Petrov, N. Craig, C. Csáki, A. X. El-Khadra, L. Baudis, J. Hall, K. T. Lesko, J. L. Orrell, J. Gonski, F. Psihas, and S. M. Simon. *Report of the 2021 U.S. Community Study on the Future of Particle Physics (Snowmass 2021) Summary Chapter*. 2023. arXiv: 2301.06581 [hep-ex].

246. A. Abada et al. “FCC-ee: The Lepton Collider”. *The European Physical Journal Special Topics* 228:2, 2019, pp. 261–623. DOI: 10.1140/epjst/e2019-900045-4. URL: <https://doi.org/10.1140/epjst/e2019-900045-4>.
247. FCC Collaboration. *FCC feasibility study status, EPS conference 2023*. Accessed: 19-12-2023. URL: https://indico.desy.de/event/34916/contributions/142198/attachments/84437/111884/230824_FCC-FS-Status-ap.pdf.
248. T. C. S. Group. *CEPC Conceptual Design Report: Volume 1 - Accelerator*. 2018. arXiv: 1809.00285 [physics.acc-ph].
249. *ILC technical design report*. Accessed: 27-10-2023. URL: <https://linearcollider.org/technical-design-report/>.
250. E. Sicking and R. Ström. “From precision physics to the energy frontier with the Compact Linear Collider”. *Nature Physics* 16:4, 2020, pp. 386–392. DOI: 10.1038/s41567-020-0834-8. URL: <https://doi.org/10.1038/s41567-020-0834-8>.
251. K. L. Bane, T. L. Barklow, M. Breidenbach, C. P. Burkhardt, E. A. Fauve, A. R. Gold, V. Heloin, Z. Li, E. A. Nanni, M. Nasr, M. Oriunno, J. M. Paterson, M. E. Peskin, T. O. Raubenheimer, and S. G. Tantawi. *An Advanced NCRF Linac Concept for a High Energy e^+e^- Linear Collider*. 2019. arXiv: 1807.10195 [physics.acc-ph].
252. B Foster, R D’Arcy, and C. A. Lindstrøm. “A hybrid, asymmetric, linear Higgs factory based on plasma-wakefield and radio-frequency acceleration”. *New Journal of Physics* 25:9, 2023, p. 093037. DOI: 10.1088/1367-2630/acf395. URL: <https://doi.org/10.1088/1367-2630/acf395>.
253. N. Bartosik, A. Bertolin, L. Buonincontri, M. Casarsa, F. Collamati, A. Ferrari, A. Ferrari, A. Gianelle, D. Lucchesi, N. Mokhov, M. Palmer, N. Pastrone, P. Sala, L. Sestini, and S. Striganov. “Detector and Physics Performance at a Muon Collider”. *Journal of Instrumentation* 15:05, 2020, P05001. DOI: 10.1088/1748-0221/15/05/P05001. URL: <https://dx.doi.org/10.1088/1748-0221/15/05/P05001>.
254. R. Barlow. “Fundamental Concepts”. In: *Data Analysis in High Energy Physics*. John Wiley & Sons, Ltd, 2013. Chap. 1, pp. 1–26. ISBN: 9783527653416. DOI: <https://doi.org/10.1002/9783527653416.ch1>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9783527653416.ch1>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/9783527653416.ch1>.
255. F. James and M. Roos. “MINUIT: a system for function minimization and analysis of the parameter errors and corrections”. *Comput. Phys. Commun.* 10:CERN-DD-75-20, 1975, pp. 343–367.
256. ATLAS Collaboration. “Measurement of the W -boson mass in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”. *Eur. Phys. J. C* 78, 2018, p. 110. DOI: 10.1140/epjc/s10052-017-5475-4. arXiv: 1701.07240 [hep-ex].

257. P.D. Group et al. “Review of Particle Physics”. *Progress of Theoretical and Experimental Physics* 2020:8, 2020, p. 083C01. ISSN: 2050-3911. DOI: 10 . 1093 / ptep / ptaa104. eprint: [https :
/ / academic . oup . com / ptep / article - pdf / 2020 / 8 / 083C01 / 34673722 / ptaa104 . pdf](https://academic.oup.com/ptep/article-pdf/2020/8/083C01/34673722/ptaa104.pdf). URL:
<https://doi.org/10.1093/ptep/ptaa104>.

SELBSTSTÄNDIGKEITSERKLÄRUNG

Ich erkläre, dass ich die Dissertation selbständig und nur unter Verwendung der von mir gemäß §7 Abs. 3 der Promotionsordnung der Mathematisch-Naturwissenschaftlichen Fakultät, veröffentlicht im Amtlichen Mitteilungsblatt der Humboldt-Universität zu Berlin Nr. 42/2018 am 11.07.2018 angegebenen Hilfsmittel angefertigt habe.