UNIVERSITY OF COPENHAGEN
FACULTY OF SCIENCE

Doctoral Thesis

Searching for Heavy Neutral Leptons at CERN

Jean-Loup Tastet

Supervisor: Oleg Ruchayskiy

February 28, 2021

This thesis has been submitted to the PhD School of The Faculty of Science, University of Copenhagen

Abstract

The Standard Model of particle physics (SM) is our best description of matter and interactions at subatomic scales. Despite its flawless record at describing the results of high-energy experiments, it cannot be a fundamental theory, for it fails to describe a number of *well-established observational phenomena*: it contains massless neutrinos (in contradiction to the observed neutrino flavor oscillations), cannot explain the observed matter-antimatter asymmetry of our Universe, and does not provide a candidate for the elusive dark matter.

One of the simplest extensions of the Standard Model which could address several — if not all — of these shortcomings consists in adding back the "missing" gauge singlet counterparts to neutrinos. These SM singlets can have a Majorana mass, whose scale is a priori unknown. If this Majorana mass is at or below the electroweak scale, the corresponding mass eigenstates — *heavy neutral leptons* (HNLs) — would interact solely through a small mixing with neutrinos. As a prime example of feebly interacting particles, they might have evaded detection so far due to their tiny interactions. HNLs are currently being actively searched by multiple experiments, and are among the main motivations for future "intensity-frontier" facilities, which will be uniquely sensitive to rare processes.

This thesis, presented as a collection of three articles, investigates phenomenological aspects of heavy neutral leptons, in relation to their search at current or proposed experiments. It concentrates on testing those properties of HNLs which are essential for resolving the aforementioned shortcomings of the SM. The first article discusses whether one could test the HNL mass degeneracy — a core requirement for HNLs to generate the observed baryon asymmetry of the Universe — by observing their oscillations at the proposed SHiP experiment. The second investigates whether a new search channel at the NA62 experiment could be used to close a currently unconstrained region in parameter space. The last article reinterprets the results of an existing experimental search for HNLs by the ATLAS experiment within a minimal yet realistic model of neutrino oscillations. By providing a scheme which allows to easily recast their exclusion limits for arbitrary model parameters, this work could greatly increase the scientific return of collider searches for HNLs.

In summary, this thesis demonstrates how a minimal, realistic model of heavy neutral leptons can nonetheless have a rich phenomenology, and discusses some important implications for experiments. In particular, it highlights the impact of model assumptions on experimental limits, and the need to interpret results within realistic models.

Synopsis

Inden for partikelfysik giver Standardmodellen (SM) os den bedste beskrivelse af stof og interaktioner ved subatomare skalaer. På trods af at SM stemmer overens med tidligere højenergieksperimenter, kan den ikke betegnes som en fundamental teori, da den ikke korrekt beskriver en række *veletablerede observerbare fænomener*: herunder indeholder den masseløse neutrinoer (hvilket modsiger de observerede neutrino-oscillationer), den forklarer ikke den observerede stof-antistof-asymmetri i vores univers, og den foreskriver ikke en kandidat til det flygtige mørke stof.

En af de simpleste udvidelser af Standardmodellen, som kunne løse en del - hvis ikke alle - af disse mangler, består i at tilføje de "manglende" højrehåndede gauge singlet modparter til neutrinoerne. Disse SM-singlets kan have en Majorana-masse, hvis skala er ukendt a priori. Hvis denne Majorana-masse er mindre end den elektrosvage skala, vil de tilsvarende masse-egentilstande — *heavy neutral leptons* (HNL'er) — udelukkende interagere ved at mixe svagt med neutrinoerne. Da de er et hovedeksempel på svagt interagerende partikler, kan dette muligvis være forklaringen på, hvorfor de indtil videre har kunnet undgå at blive detekteret. På nuværende tidspunkt søges HNL'er aktivt efter ved flere eksperimenter, og de er blandt hovedmålene for fremtidige "højintensitetseksperimenter", som vil have en unik følsomhed over for disse sjældne processer.

Denne afhandling, der præsenteres som en samling af tre artikler, undersøger fænomenologiske aspekter af heavy neutral leptons i forbindelse med, hvordan de bliver søgt efter ved nuværende eller fremtidige eksperimenter. Den sporer ind på at teste egenskaberne for HNL'er, som er essentielle for at finde løsninger på de førnævnte mangler i SM. Den første artikel diskuterer, hvorvidt man kan undersøge massedegenerationen af HNL'er — som kræves, for at HNL'er genererer baryon-asymmetrien i universet — ved at observere deres oscillationer ved det fremtidige SHiP-eksperiment. Den anden artikel undersøger, hvorvidt en ny kanal ved NA62-eksperimentet kunne bruges til at afgrænse det ellers ubegrænsede parameterrum. Den sidste artikel genfortolker resultaterne af en eksperimentel søgen efter HNL'er ved ATLAS-eksperimentet ved brug af en minimal, dog realistisk model for neutrino-oscillationer. Ved at give en metode, som gør det muligt enkelt at genfortolke afgrænsingen for vilkårlige modelparametre, kan dette arbejde i høj grad øge den videnskablige værdi af søgninger efter HNL'er ved collider-eksperimenter.

For at opsummere, demonstrerer denne afhandling, hvordan en minimal, realistisk model for heavy neutral leptons kan have en rig fænomenologi, og den diskuterer nogle vigtige følger for eksperimenterne. Den understreger i særdeleshed, hvilken virkning antagelser for modeller har på eksperimentelle afgrænsninger samt nødvendigheden af at fortolke resultater med realistiske modeller.

Acknowledgements

Despite being one of the very few phenomenologists in the experimental particle physics group, I have received help and support from many people during my three years at the NBI. Their contribution to my PhD journey shall not be understated, and I would like to express my gratitude to them.

First and foremost, I would like to thank my supervisor, Oleg, for the trust he placed in me by hiring me as his first PhD student, and for allowing me to follow some of my ideas while helping me get back on track whenever I hit a dead end. He helped instill a healthy dose of pragmatism in me, by pushing me to publish my papers before the heat death of the Universe, instead of starting to check my results all over again for the $(n+1)^{\text{th}}$ time. I very much appreciated his openness and his ability to listen, which have led to many constructive discussions; this was particularly critical when I had to quickly adapt my project as my experiment was almost cancelled at the same time as the world was hit by a major pandemic.

I also want to thank Inar for our lively and never-ending discussions on Telegram about physics and recent papers, from which I have learnt a lot, and which have kept me from getting bored and allowed me to remain up to date with the latest developments in our relatively small field. I am also particularly thankful to him for supervising me in the early stages of this PhD, and then for providing guidance when needed and for sharing with me his knowledge of cosmology and leptogenesis.

This also extends to my former group at EPFL, for their occasional but important contributions to some of our discussions. In particular, I would like to thank Misha for his always very thoughtful suggestions, as well as Sonia, to whom Inar promised that I was an expert in particle kinematics, prompting me to start learning about this topic which, despite all appearances, turned out to be fun and lead to some interesting results.

I also want to thank all those with whom I have had the chance to collaborate: Evgueni, Kyrylo, Maksym; as well as my ATLAS colleagues who gave me the unique opportunity to take an active part in HNL searches, and in particular Stefania, Monika, Dominique and Alison; and finally all those who embarked on the crazy project of building a new experiment dedicated to the search for particles whose defining property is to be very well hidden. This extends to the people who took the lead in developing the emerging field of feebly interacting particles, and managed to create a small but lively community around it.

Finally, I want to thank my PhD, postdoc and master comrades and — more importantly — friends, for welcoming a theorist (and a French one, on top of that) in their group, and for all the amazing discussions we had during breaks and in Søernes (or, in 2020 fashion, during virtuøl). Christian, Emil, Emilie, Étienne, Fabian, Freja, Helene, Jake, Katka, Laís, Meera, Vojta, Vytautas, Yuko, Zlatko, Zuzana: you deserve all the credit for my survival during these harsh Danish winters, and for keeping me (more or less) sane during the past year! I look forward to finally being part of the very select "duck club". This also extends to all the people whom I used to occasionally meet at the NBI, and who contributed to making it a lively place. Special thanks to Julie for helping me translate the abstract in Danish, a language which far exceeds the learning capabilities of my French brain.

I would also like to thank my family and my friends outside Denmark, for keeping me connected with the outside world through uncountable "skypéros". I am especially grateful to my parents, who did not hesitate to travel through Europe during a pandemic in order to help me move (twice!) while I was finalizing the present thesis.

Contents

Chapter 1

Introduction

Contents

1.1 The Standard Model of particle physics

1.1.1 Brief introduction[1]

The Standard Model of particle physics [3–6] (SM) is the culmination of more than a century of experimenting with matter at the subatomic scale. It is the quantum field theory (QFT) which, as of today, best describes the known forms of matter and their interactions (with the notable exception of gravity). While quantum electrodynamics only merged quantum mechanics with electrodynamics, the Standard Model generalizes the approach to the weak and strong forces by describing them as non-abelian (Yang-Mills) gauge theories. More specifically, the Standard Model is based on the gauge group $SU(3)_\mathrm{c} \times SU(2)_\mathrm{L} \times U(1)_\mathrm{Y}$ which defines the structure of its various interactions. Their strength, as well

[1]This brief introduction is inspired by ref. [1], but focuses on the aspects which are most relevant to this thesis. In order to simplify the description of Majorana particles, fermions are described using two-component spinors, following the formalism from ref. [2].

as the particle content of the SM, are only fully specified by its *Lagrangian* $\mathcal{L}_{\text{SM}}$, which contains a number of terms.

First, the gauge term $\mathcal{L}_{\text{gauge}}$ defines the structure of interactions and the strength of their self-interaction:

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4}\sum_{i=1...8} G^{i\mu\nu}G^{i}{}_{\mu\nu} - \frac{1}{4}\sum_{j=1,2,3} W^{j\mu\nu}W^{j}{}_{\mu\nu} - \frac{1}{4}B^{\mu\nu}B_{\mu\nu} \tag{1.1}$$

where $G^{i\mu\nu}$, $W^{j\mu\nu}$ and $B^{\mu\nu}$ respectively denote the gauge fields in the adjoint representations of $SU(3)_{\text{c}}$, $SU(2)_{\text{L}}$ and $U(1)_{\text{Y}}$, defined as:

$$\begin{aligned} G^{i\mu\nu} &= \partial^\mu G^{i\nu} - \partial^\nu G^{i\mu} - g_s f_{ijk} G^{j\mu} G^{k\nu} \\ W^{i\mu\nu} &= \partial^\mu W^{i\nu} - \partial^\nu W^{i\mu} - g \varepsilon_{ijk} W^{j\mu} W^{k\nu} \\ B^{\mu\nu} &= \partial^\mu B^{\nu} - \partial^\nu B^{\mu} \end{aligned}$$

with $G^{i\mu}$, $W^{i\mu}$ and B^{μ} respectively denoting the fundamental representations of these same fields, f_{ijk} and ε_{ijk} the respective structure constants of $SU(3)$ and $SU(2)$, and g_s and g the respective couplings constants of the strong and weak forces.

The matter content of the Standard Model is defined by the fermionic term:

$$\mathcal{L}_f = \sum_{\alpha=1,2,3} \left(\sum_{f=Q_{\text{L}},L_{\text{L}}} i f^\dagger_\alpha \bar{\sigma}^\mu D_\mu f_\alpha + \sum_{f^\dagger = u^\dagger_{\text{R}}, d^\dagger_{\text{R}}, e^\dagger_{\text{R}}} i f_\alpha \sigma^\mu D_\mu f^\dagger_\alpha \right) \tag{1.2}$$

where α is the generation / flavor index, f denotes the left-handed fermion fields and $f^\dagger$ the right-handed[2,3] ones (following the conventions from ref. [2], with the fermion indices omitted), $\bar{\sigma}^\mu$ and σ^μ are the "covariant" Pauli matrices, and D_μ is the covariant derivative:

$$D_\mu = \partial_\mu + i g_s \frac{\lambda^i}{2} G^i_\mu + i g \frac{\sigma^j}{2} W^j_\mu + i g' q_{\text{Y}} B_\mu \tag{1.3}$$

with, respectively, g_s, g and g' the coupling constants for $SU(3)_{\text{c}}$, $SU(2)_{\text{L}}$ and $U(1)_{\text{Y}}$, i and j the adjoint representation indices for $SU(3)_{\text{c}}$ and $SU(2)_{\text{L}}$, λ^i the Gell-Mann matrices, σ^j the Pauli matrices (those are the generators of the respective Lie algebras, up to a factor of 1/2), and q_{Y} the hypercharge. In eq. (1.2), the "left-handed" fields $Q_{\text{L}} = (u_{\text{L}}, d_{\text{L}})$ and $L_{\text{L}} = (\nu_{\text{L}}, e_{\text{L}})$ are $SU(2)_{\text{L}}$ doublets while the "right-handed" ones $u^\dagger_{\text{R}}$, $d^\dagger_{\text{R}}$ and $e^\dagger_{\text{R}}$ are singlets. Note the absence of a right-handed ($SU(2)_{\text{L}}$ singlet) neutrino field $\nu^\dagger_{\text{R}}$ in this "classical" version of the Standard Model.

[2] Note that there are a number of close but different definitions of "handedness". Here, by "left-handed" or "right-handed" we will never refer to the helicity of a specific particle; rather, we use it to denote the chirality (i.e. the corresponding spin-$\frac{1}{2}$ representation) of the field. Since the hermitian conjugate of a field with left chirality is equivalent to a field with right chirality (and reciprocally) — as our notation makes manifest — to avoid confusion we will denote by "left-handed" and use the $_{\text{L}}$ subscript for the $SU(2)_{\text{L}}$ doublets (for which the "particle" field is left-handed, and the anti-particle field right-handed) and by "right-handed" (with the $_{\text{R}}$ subscript) the $SU(2)_{\text{L}}$ singlets (for which the opposite is true).

[3] Note that $f\sigma^\mu \partial_\mu f^\dagger = f^\dagger \bar{\sigma}^\mu \partial_\mu f + \partial_\mu (f \sigma^\mu f^\dagger)$ [7], where the last term is a total divergence that is usually irrelevant.

These first two terms are sufficient to define the gauge structure of the SM and the interactions of fermions at sufficiently high energies. But they still lack a crucial ingredient: masses, which are necessary to account for the observed fermion masses as well as for the "weakness" of the weak interaction at low energies. Masses have proven to be particularly tricky to implement in gauge theories without breaking the gauge invariance. In the Standard Model, it is the well-known Higgs mechanism [8–11] which generates the masses of fermions and bosons while preserving gauge invariance. It introduces a new pair of complex scalar fields $\phi = (\phi^+, \phi^0)$ — the Higgs field — which transforms as an $SU(2)_\mathrm{L}$ doublet with hypercharge $y_\phi = \frac{1}{2}$ and couples to left-right fermion pairs (of potentially different generations) through the so-called Yukawa couplings $Y^f_{\alpha\beta}$, $f = u, d, e$:

$$\mathcal{L}_\mathrm{Yuk} = -\sum_{\alpha,\beta=1,2,3} \left(Y^u_{\alpha\beta} (Q^\dagger_{\mathrm{L}\alpha} \cdot \tilde{\phi}) u^\dagger_{\mathrm{R}\beta} + Y^d_{\alpha\beta} (Q^\dagger_{\mathrm{L}\alpha} \cdot \phi) d^\dagger_{\mathrm{R}\beta} + Y^e_{\alpha\beta} (L^\dagger_{\mathrm{L}\alpha} \cdot \phi) e^\dagger_{\mathrm{R}\beta} \right) + \text{h.c.} \tag{1.4}$$

where $\tilde{\phi} = \varepsilon \phi^\dagger$. Note the absence of a Yukawa term for neutrinos, which lack a right-handed / $SU(2)_\mathrm{L}$ singlet component in the SM.

In the Standard Model, the dynamics of the Higgs field are described using an ad-hoc potential $V(\phi)$, built by imposing the requirement of $SU(2)_\mathrm{L} \times U(1)_\mathrm{Y}$ gauge invariance, keeping only the relevant terms of dimension ≤ 4 and requiring the presence of a non-trivial minimum at $\phi \neq 0$:

$$V(\phi) = \mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2 \quad \text{where } \mu^2 < 0 \text{ and } \lambda > 0 \tag{1.5}$$

resulting in the Lagrangian term:

$$\mathcal{L}_\phi = (D^\mu \phi)^\dagger (D_\mu \phi) - V(\phi) \tag{1.6}$$

with the usual covariant derivative $D_\mu \phi = (\partial_\mu + i\frac{g}{2}\sigma^j W^j_\mu + i\frac{g'}{2} B_\mu)\phi$. This formally completes the Standard Model Lagrangian:

$$\mathcal{L}_\mathrm{SM} = \mathcal{L}_\mathrm{gauge} + \mathcal{L}_f + \mathcal{L}_\mathrm{Yuk} + \mathcal{L}_\phi \tag{1.7}$$

However, due to the non-trivial minimum of the potential at $|\phi| = \sqrt{-\mu^2/2\lambda}$, the ground state (or vacuum) of the theory corresponds to a non-zero *vacuum expectation value* (v.e.v) $\langle\phi\rangle$ of the Higgs doublet. Although the $SU(2)_\mathrm{L} \times U(1)_\mathrm{Y}$ invariance allows to freely choose this value using a suitable gauge transformation, the ground state of the theory is no longer invariant under this gauge group once the value is set (i.e. the choice of the vacuum acts as a gauge fixing condition). A typical choice for representing the Higgs field in the broken phase is the so-called unitary gauge:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix} \tag{1.8}$$

where $v = \sqrt{-\mu^2/\lambda}$ and h is the dynamical part of the Higgs field in the broken phase, i.e. the Higgs boson, first observed by ATLAS [12] and CMS [13] in 2012 at the Large Hadron Collider (LHC) at CERN, with a (tree-level) mass $m_h = \sqrt{-2\mu^2} \approx 125\,\mathrm{GeV}$. Even in this broken phase, the Higgs field retains a residual $U(1)$ gauge invariance, generated by $Q = \frac{1}{2}\sigma^3 + Y$, which is nothing else than the electromagnetic charge we are all familiar with. This phenomenon (of the Higgs field acquiring a non-trivial v.e.v) is called the *electroweak symmetry breaking* (EWSB), and is often expressed as

$$SU(2)_{\mathrm{L}} \times U(1)_{\mathrm{Y}} \xrightarrow{\text{EWSB}} U(1)_{\mathrm{Q}}.$$

The two main consequences of the Higgs mechanism come from the Yukawa couplings and from the Higgs kinetic term. Of most interest to this thesis is the Yukawa term. After expanding the Higgs field (1.8) in eq. (1.4), we obtain Dirac mass terms[4] for left-right fermion pairs:

$$\mathcal{L}_{\mathrm{Yuk}} \xrightarrow{\text{EWSB}} -\sum_{\alpha,\beta=1,2,3}\left(m^u_{\alpha\beta} u_{\mathrm{L}\alpha} u_{\mathrm{R}\beta} + m^d_{\alpha\beta} d_{\mathrm{L}\alpha} d_{\mathrm{R}\beta} + m^e_{\alpha\beta} e_{\mathrm{L}\alpha} e_{\mathrm{R}\beta}\right) + \text{h.c.} + \mathcal{O}(h) \tag{1.9}$$

where the Dirac masses are obtained as $m^f_{\alpha\beta} = \frac{v}{\sqrt{2}}(Y^f_{\alpha\beta})^*$. These mass "matrices" (with flavor indices $1, 2, 3$ corresponding to the three generations) are in general non-diagonal. As arbitrary complex 3×3 matrices, they can only be diagonalized by applying different unitary transformations $V^f_{\mathrm{L}\alpha i}$ and $V^f_{\mathrm{R}\alpha i}$ to the left- and right-handed fields:[5]

$$f_{\mathrm{L}\alpha} = V^f_{\mathrm{L}\alpha i} f'_{\mathrm{L}i} \tag{1.10}$$

$$f_{\mathrm{R}\beta} = V^f_{\mathrm{R}\beta i} f'_{\mathrm{R}i} \tag{1.11}$$

$$m^f_{\alpha\beta} f_{\mathrm{L}\alpha} f_{\mathrm{R}\beta} = (\underbrace{V^f_{\mathrm{L}\alpha i} V^f_{\mathrm{R}\beta j} m^f_{\alpha\beta}}_{=m^f_i \delta_{ij}}) f'_{\mathrm{L}i} f'_{\mathrm{R}j} = m^f_i f'_{\mathrm{L}i} f'_{\mathrm{R}i} \tag{1.12}$$

(with implicit summation over repeated indices, and $'$ indicating fields in the diagonal basis). These unitary transformations leave all kinetic terms invariant:

$$i f^\dagger_{\mathrm{L}\alpha} \bar{\sigma}^\mu D_\mu f_{\mathrm{L}\alpha} = i(V^f_{\mathrm{L}\alpha i} f'_{\mathrm{L}i})^\dagger \bar{\sigma}^\mu D_\mu V^f_{\mathrm{L}\alpha j} f'_{\mathrm{L}j} = i f'^\dagger_{\mathrm{L}i} \bar{\sigma}^\mu D_\mu f'_{\mathrm{L}i} \tag{1.13}$$

$$i f_{\mathrm{R}\alpha} \sigma^\mu D_\mu f^\dagger_{\mathrm{R}\alpha} = i V^f_{\mathrm{R}\alpha i} f'_{\mathrm{R}i} \sigma^\mu D_\mu (V^f_{\mathrm{R}\alpha j} f'_{\mathrm{R}j})^\dagger = i f'_{\mathrm{R}i} \sigma^\mu D_\mu f'^\dagger_{\mathrm{R}i} \tag{1.14}$$

For the right-handed (singlet) fields, one can always perform such a transformation. However, for the left-handed (doublet) ones, the same transformation must be applied to the entire $SU(2)_{\mathrm{L}}$ doublet in order not to spoil the gauge invariance of the theory. Therefore the mass matrices of (left-handed) fields from a same doublet (such as u_{L} and d_{L}) can never be *simultaneously* diagonalized. E.g. if we choose to diagonalize the Dirac mass matrix of "up" quarks, then the "dynamical" down quarks (which enter the kinetic term and couple to vector bosons through the covariant derivative) will be different from the

[4]Note that Majorana mass terms, of the form $mff + \text{h.c.}$ (i.e. coupling a fermion field f with itself) would require Yukawa terms which are not compatible with the gauge symmetries of this model. Such Majorana mass terms are only possible for Standard Model singlets.

[5]Our choice of parametrization (using two-component spinors [2] f_{L} and f_{R}) makes it manifest that the left- and right-handed fields are separate degrees of freedom which can be operated on independently.

"massive" down quarks (which are the propagating states, in the limit of asymptotic freedom). The mismatch between the dynamical and massive quarks is quantified by the Cabibbo-Kobayashi-Maskawa (CKM) matrix:

$$(V_{\mathrm{CKM}})_{\alpha\beta} = (V^u_{\mathrm{L}})_{\alpha i}(V^{d\dagger}_{\mathrm{L}})_{i\beta} = V^u_{\mathrm{L}\alpha i}V^{d*}_{\mathrm{L}\beta i} \tag{1.15}$$

This phenomenon is referred to as *mixing*. A similar phenomenon, in the case of neutrinos, will be of crucial importance for this thesis.

The second consequence of the Higgs mechanism is the generation of (gauge invariant) mass terms for three electroweak vector bosons. Due to its charge under $SU(2)_{\mathrm{L}} \times U(1)_{\mathrm{Y}}$, the Higgs field couples to all the electroweak vector bosons through the covariant derivative $D_\mu\phi = (\partial_\mu + i\frac{g}{2}\sigma^j W^j_\mu + i\frac{g'}{2}B_\mu)\phi$. After electroweak symmetry breaking, and defining $W^\pm_\mu = \frac{1}{\sqrt{2}}(W^1_\mu \mp iW^2_\mu)$, the kinetic term $(D^\mu\phi)^\dagger(D_\mu\phi)$ generates a mass term for the electroweak vector bosons:

$$(D^\mu\phi)^\dagger(D_\mu\phi) \xrightarrow{\mathrm{EWSB}} M_W^2 W^{+\mu}W^-_\mu + \frac{M_Z^2}{2}Z^\mu Z_\mu + \mathcal{O}(h) + \mathcal{O}(\partial h) \tag{1.16}$$

where

$$M_W = \frac{gv}{2} \qquad M_Z = \frac{\sqrt{g^2+g'^2}v}{2} \qquad Z_\mu = \frac{-g'B_\mu + gW^3_\mu}{\sqrt{g^2+g'^2}}$$

These vector bosons have absorbed the three "broken" components of the Higgs doublet by each acquiring a new longitudinal polarization. Defining $A_\mu = (g'W^3_\mu + gB_\mu)/\sqrt{g^2+g'^2}$, we see that this combination of fields produces a zero contribution to the covariant derivative $D_\mu\phi$ in the unitary gauge. Therefore this last vector field does not get a mass from the Higgs mechanism, nor does it couple to it. It is nothing else than the photon field of electrodynamics.

The model described above is still, at the time of writing, our best description of nuclear and high-energy physics. Unchallenged by accelerator experiments, which have largely contributed to its development, it nonetheless fails decisively to describe a number of well-established observations. These are the topic of the next section.

1.1.2 Observational shortcomings of the Standard Model

Neutrino masses

Starting with the Homestake experiment [14] in 1968, it was observed that the flux of electron neutrinos from the Sun did not match the expectation from the Standard Model (in which neutrino flavor is conserved). This "solar neutrino problem" lasted for more than three decades, until Super-Kamiokande [15] in 1998 and SNO [16] in 2001 finally produced decisive evidence for lepton flavor non-conservation. By looking at the atmospheric neutrino flux as a function of the zenith angle (and thus of the distance travelled through the Earth) and observing an angle dependent deficit of muon neutrinos, Super-Kamiokande provided conclusive evidence for neutrino oscillations, first proposed in refs. [17, 18]. It is SNO which resolved the actual solar neutrino problem, by measuring both the electron neutrino flux (through charged-current interactions) as well as the rate of elastic scattering of electrons (which is sensitive to the fluxes of all three flavors). By observing a ratio of electron flavor < 1, it gave credence to the hypothesis that neutrinos produced inside the Sun undergo resonant flavor transitions due to the MSW effect [19, 20]. Both neutrino oscillations and the MSW effect can only be explained if neutrinos have a (very small) mass, in direct contradiction with the usual formulation of the Standard Model, in which they are massless.

Neutrino oscillations can be schematically understood using a simple quantum mechanical argument[6]. If neutrinos are massive, there will be a mismatch between the mass eigenbases of charged leptons and neutrinos, like in the quark sector. This mismatch is quantified by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix V_{PMNS} [21]. Neutrinos are typically produced in charged-current decays, where they couple to a charged lepton. Unlike neutrinos, charged leptons have very distinct masses, and thus constitute distinguishable quantum states: $|e\rangle$, $|\mu\rangle$ and $|\tau\rangle$. It is therefore convenient to use them for defining the corresponding neutrino flavor (i.e. interacting) states $|\nu_e\rangle$, $|\nu_\mu\rangle$ and $|\nu_\tau\rangle$. They are related to the massive neutrinos states[7] $|\nu_i\rangle$ ($i = 1, 2, 3$) by:

$$|\nu_\alpha\rangle = (V_{\text{PMNS}})^*_{\alpha i} |\nu_i\rangle \tag{1.17}$$

Now consider the space-time evolution of a flavour state $|\nu_\alpha(x)\rangle$, within the Schrödinger picture. This is most conveniently done by applying the space-time translation operator $e^{-i\hat{P}\cdot x}$ to the quantum state:

$$|\nu_\alpha(x)\rangle = e^{-i\hat{P}\cdot x} |\nu_\alpha\rangle = (V_{\text{PMNS}})^*_{\alpha i} e^{-i\hat{P}\cdot x} |\nu_i\rangle = (V_{\text{PMNS}})^*_{\alpha i} e^{-ip_i\cdot x} |\nu_i\rangle \tag{1.18}$$

where p_i is the momentum of the plane wave described by the state $|\nu_i\rangle$. We now need to make two assumptions which can only be justified with a rigorous quantum field theoretical treatment. We assume $x = (t, \mathbf{x})$ to be the same for all the mass eigenstates[8], and in addition we also assume the three-momentum $\mathbf{p}$ to be the same[9], such that $p_i = (E_i, \mathbf{p})$ with $E_i = \sqrt{m_i^2 + \mathbf{p}^2}$.

After a finite time, the flavor state $|\nu_\alpha\rangle$ will have evolved into a state $|\nu_\alpha(x)\rangle$ with a non-zero overlap with the other flavor states $|\nu_\beta\rangle$: it has *oscillated*. The probability of such a transition can be computed using Born's rule:

$$\begin{aligned} P_{\alpha\to\beta}(x) = \left|\langle\nu_\beta|\nu_\alpha(x)\rangle\right|^2 &= \Big|(V_{\text{PMNS}})_{\beta j}(V_{\text{PMNS}})^*_{\alpha i} e^{-ip_i\cdot x} \underbrace{\langle\nu_j|\nu_i\rangle}_{=\delta_{ij}}\Big|^2 \\ &= \left|(V_{\text{PMNS}})_{\beta i}(V_{\text{PMNS}})^*_{\alpha i} e^{-ip_i\cdot x}\right|^2 \end{aligned} \tag{1.19}$$

For an ultra-relativistic neutrino measured at a distance L (in the laboratory) from its production site, the phase can be well approximated as:

$$p_i \cdot x = E_i t - |\mathbf{p}| L \cong (E_i - |\mathbf{p}|) L \cong \frac{m_i^2}{2|\mathbf{p}|} L \tag{1.20}$$

[6] A much more accurate model will be used in section 1.3.2 to describe HNL oscillations. This model is also suitable for describing the oscillations of light neutrinos.

[7] The relation $\nu_{\text{L}\alpha} = (V_{\text{PMNS}})_{\alpha i}\nu_i$ for fields implies $|\nu_\alpha\rangle = (V_{\text{PMNS}})^*_{\alpha i} |\nu_i\rangle$ for states.

[8] In the external wave packet model, this is justified because the interference is computed by integrating the wave packet product $\psi_\beta(x)^*\psi_\alpha(x)$ over the same space-time coordinate x.

[9] In a QFT treatment, this is justified because we can choose $\mathbf{p}$ as the common integration variable when integrating over phase space.

leading to the well-known formula for the oscillation probability:

$$P_{\alpha\to\beta}(L) = \left|(V_{\mathrm{PMNS}})^*_{\alpha i}(V_{\mathrm{PMNS}})_{\beta i} e^{-i\frac{m_i^2}{2|\mathbf{p}|}L}\right|^2 \tag{1.21}$$

This simplified quantum mechanical treatment has a number of issues, which are discussed at length in ref. [22] and references therein. When we will later describe the oscillations of heavy neutral leptons, we will therefore use a more robust quantum field theoretical treatment.

Baryon asymmetry of the Universe

The observable Universe contains a manifest excess of matter over anti-matter. This *baryon asymmetry of the Universe* (BAU) can be quantified by the ratio η of the baryon number density over the photon density:

$$\eta = \frac{n_b - n_{\bar{b}}}{n_\gamma} \sim 6 \cdot 10^{-10} \tag{1.22}$$

where n_b, $n_{\bar{b}}$ and n_γ respectively denote the number densities of baryons, anti-baryons and photons.

In order for this matter-antimatter asymmetry to be dynamically generated (by opposition to resulting from initial conditions or from a local overdensity of matter which must be at least of the size of the observable universe [23]), it has been shown that three conditions — the so-called *Sakharov conditions* [24] — must be satisfied:

- *The baryon number B cannot be conserved.* Although B is conserved at the perturbative level in the SM, this condition is actually satisfied, since the non-perturbative "sphaleron" processes [25] violate $B + L$ (but preserve $B - L$), where L stands for the total lepton number.
- *CP symmetry (and therefore T symmetry) must be violated.* This is to ensure that the baryon asymmetry generated by a given B-violating process is not compensated by its CP conjugate. Although the Standard Model contains CP violation in the quark sector due to the CKM matrix, it turns out to be too small to explain the observed value of η [23].
- *At some point during its evolution, the Universe must depart from thermodynamic equilibrium.* This is because, at equilibrium, the forward and reverse processes compensate each other exactly, such that all densities remain constant.

These three conditions are necessary, but not sufficient, to generate the observed baryon asymmetry of the Universe. In particular, the Standard Model cannot generate the observed value of η despite satisfying these conditions [23]. It therefore needs to be extended. The mechanisms capable of dynamically generating a sufficiently large baryon asymmetry are collectively referred to as *baryogenesis*.

Dark matter

A number of cosmological and astrophysical observations — such as galaxy rotation curves, peculiar velocities of galaxies in clusters, the age of galaxies and cosmological fits (see ref. [26], ch. 27 for a review) — cannot be explained solely with the Standard Model (along with general relativity). Among the proposed solutions to these problems, the most consensual — and arguably the most parsimonious, in the sense that it can resolve all of these problems simultaneously — consists in postulating the existence of a new, invisible form of matter: *dark matter*. The nature of dark matter is essentially unconstrained; only its abundance is known. However, in the spirit of the Standard Model, it is not unreasonable to expect dark matter to take the form of one or more new particles, which would couple only very feebly to the SM in order to remain "dark" (i.e. not interact electromagnetically) and stable on cosmological time scales. The possibility that such *feebly interacting particles* exist beyond the SM and may have evaded detection so far is currently being taken very seriously, and constitutes the primary physics goal of a new generation of experiments. Among these particles are *heavy neutral leptons*, which will be the main focus of this thesis.

1.2 Heavy Neutral Leptons

1.2.1 Motivation

The observations of the MSW effect and of neutrino flavor oscillations (discussed in section 1.1.2) provide unambiguous evidence for physics beyond the Standard Model. Indeed, both of these effects rely on non-zero mass splittings between different mass eigenstates, implying that at least two neutrino mass eigenstates must have non-zero masses. This is in direct contradiction with the SM, in which neutrinos are massless Weyl fermions. In order to give neutrinos a mass, it is necessary[10] to introduce additional degrees of freedom in the SM. Arguably, the simplest implementation of massive neutrinos consists in generating Dirac neutrino masses from the Yukawa couplings through the Higgs mechanism, just like for the other fermions. This requires introducing three new "right-handed" (i.e. $SU(2)_{\mathrm{L}}$ singlet) neutrinos $\nu_{\mathrm{R}\alpha}$ ($\alpha = 1, 2, 3$), effectively restoring some sort of symmetry between the quark and lepton sectors:

$$\begin{aligned}\mathcal{L}_{\mathrm{Dirac}\ \nu} &= \mathcal{L}_{\mathrm{SM}} + i\nu_{\mathrm{R}\alpha}\sigma^{\mu}\partial_{\mu}\nu_{\mathrm{R}\alpha}^{\dagger} - Y^{\nu}_{\alpha\beta}(L^{\dagger}_{\mathrm{L}\alpha}\cdot\tilde{\phi})\nu^{\dagger}_{\mathrm{R}\beta} - (Y^{\nu}_{\alpha\beta})^{*}(L_{\mathrm{L}\alpha}\cdot\tilde{\phi})\nu_{\mathrm{R}\beta} \\ &\xrightarrow{\mathrm{EWSB}} \mathcal{L}_{\mathrm{SM}} + i\nu_{\mathrm{R}\alpha}\sigma^{\mu}\partial_{\mu}\nu^{\dagger}_{\mathrm{R}\alpha} - \left(m^{\nu}_{\alpha\beta}\nu_{\mathrm{L}\alpha}\nu_{\mathrm{R}\beta} + \mathrm{h.c.}\right)\end{aligned} \tag{1.23}$$

where $m^{\nu}_{\alpha\beta} = \frac{v}{\sqrt{2}}(Y^{\nu}_{\alpha\beta})^{*}$ is the Dirac mass matrix of neutrinos and we have replaced $D_{\mu} \to \partial_{\mu}$ because these right-handed neutrinos turn out to be SM singlets, i.e. they are neutral under *all* the SM gauge interactions. One can then, similarly to the case of quarks (discussed in section 1.1.1), use a pair of unitary transformations to diagonalize *one* of the leptonic mass matrices (either $m^{e}_{\alpha\beta}$ or $m^{\nu}_{\alpha\beta}$). The usual choice is to diagonalize the charged fermion mass matrix $m^{e}_{\alpha\beta}$, since the electron, muon and tau "mass eigenstates" have very distinct signatures, which makes them more convenient than neutrinos for defining lepton "flavor". The mismatch between the interacting and massive

[10] Since the left-handed SM neutrinos are not neutral, they cannot have a Majorana mass term.

neutrino fields is then quantified by the PMNS matrix [21] $(V_{\text{PMNS}})_{\alpha\beta}$, which is the leptonic counterpart to the CKM matrix.

Although this minimal model is sufficient to account for neutrino masses and oscillations, it is often regarded as unsatisfactory, for two main reasons. First, the absolute scale of neutrino masses is quite puzzling, since it is at least 11 orders of magnitude lower than the electroweak scale [26]. There is no shortage of models to explain this oddity, which is often considered as a *hint* that neutrino masses are not produced by the same mechanism that generates the other fermion masses. Second, the new Standard Model singlets $\nu_{\text{R}\alpha}$ can also, by virtue of being completely neutral, accommodate a *Majorana* mass term:

$$\mathcal{L}_{\text{Majorana mass}} = -\frac{M_{\alpha\beta}}{2}\nu_{\text{R}\alpha}\nu_{\text{R}\beta} + \text{h.c.} \tag{1.24}$$

This term has dimension 3 and is therefore relevant. Since the Standard Model seems to contain all relevant ($d < 4$) and marginal ($d = 4$) terms which are compatible with its symmetries and can be built out of its field content, one could reasonably expect a Majorana mass term to be present for right-handed neutrinos if they are present in the model.

Among the models of neutrino mass generation, the most popular has long been the so-called *type-I seesaw mechanism* [27–34], which postulates that neutrinos acquire their mass by mixing with very heavy Majorana neutrinos through "natural" Yukawa couplings $|Y^\nu_{\alpha\beta}| \sim 1$. If the Majorana mass of the singlets $\nu_{\text{R}\alpha}$ is around some GUT scale ($|M_{\alpha\beta}| \sim 10^{15}\,\text{GeV}$), then the light neutrinos acquire at tree level a mass $m_{\text{light}} \sim {}^{\langle|\phi|\rangle^2}/_{|M_{\alpha\beta}|} \sim 0.1\,\text{eV}$, in agreement with observations. This model will be discussed in greater details in section 1.2.2.

It was later realized [35] that a type-I seesaw with a Majorana mass well below the GUT or Planck scale, combined with small (but not unnaturally so[11]) Yukawa couplings, can provide an equally viable explanation for the smallness of neutrino masses. Although there is no unambiguous prescription to determine the scale of the Majorana mass, "low-scale" seesaw models with Majorana masses at or below the electroweak scale — such as for instance the *neutrino Minimal Standard Model* [35,36] (νMSM), which posits that at least two right-handed neutrinos are nearly degenerate in mass — have received considerable attention in recent years, for two main reasons. First, the νMSM has been numerically proven to be capable of generating the observed baryon asymmetry of the universe (BAU) in a large fraction of its allowed parameter space [35–37], and in part of this parameter space the lightest right-handed neutrino even becomes a viable dark matter candidate [38]. Thus the νMSM is capable of simultaneously explaining not only neutrino masses, but also the BAU and dark matter, a feat which the Standard Model is unable to achieve. Second, the mass range of the right-handed neutrinos of the νMSM is within the reach of current facilities, the sensitivity of which is only limited by the potentially small mixing between the flavor eigenstates and the "heavy" mass eigenstates — dubbed *heavy neutral leptons* (or HNLs) owing to their experimental signature. This makes HNLs a prime target for future intensity frontier experiments, and several of them have already published sensitivity estimates [39–44].

[11] The electron Yukawa coupling in the SM (in the diagonal basis) is *measured* to be $|Y^e_{11}| \sim 3 \cdot 10^{-6}$, and a low-scale seesaw could work for $|Y^\nu_{\alpha\beta}| \gtrsim 10^{-6}$.

In order to generate a sufficiently large baryon asymmetry, the νMSM relies on the ARS mechanism [45], or *baryogenesis through sterile neutrino oscillations*, which postulates that CP-violating oscillations between at least two HNLs generate nonzero individual lepton numbers $L_{e,\mu,\tau} \neq 0$, only part of which is converted to baryon number through sphaleron transitions, due to differences in couplings U_α^2 and equilibration rates between generations. Such oscillations obviously require at least two heavy neutrino mass eigenstates to be nearly-degenerate (see [37] for accurate bounds on the degeneracy). Moreover, if the lightest HNL N_1 is sufficiently long-lived to be a viable dark matter candidate, it has been shown [38] that the νMSM can account for the observed (warm) dark matter abundance, provided that the physical mass splitting δm^2 between the two heaviest HNLs $N_{2,3}$ is much smaller than $\Delta m^2_{\mathrm{atm}}$. Such a small mass splitting could lead to an oscillation length long enough to be resolved in laboratory experiments.

1.2.2 The type-I seesaw mechanism[12]

The type-I seesaw mechanism [27–34] extends the Standard Model with $\mathcal{N}$ right-handed neutrinos $\nu_{\mathrm{R}I}$ ($I = 1 \dots \mathcal{N}$), which are spin-$\frac{1}{2}$ SM singlets with Majorana masses[13] M_I, and new Yukawa couplings $Y^\nu_{\alpha I}$:

$$\mathcal{L}_{\mathrm{seesaw}} = \mathcal{L}_{\mathrm{SM}} + \frac{i}{2}\nu^\dagger_{\mathrm{R}I}\bar{\sigma}^\mu\partial_\mu\nu_{\mathrm{R}I} - (Y^\nu_{\alpha I})^*(L_{\mathrm{L}\alpha}\cdot\tilde{\phi}^\dagger)\nu_{\mathrm{R}I} - \frac{M_I}{2}\nu_{\mathrm{R}I}\nu_{\mathrm{R}I} + \mathrm{h.c.} \tag{1.25}$$

After electroweak symmetry breaking, the Yukawa interaction generates a Dirac mass term $(m_D)_{\alpha I} = \frac{v}{\sqrt{2}}(Y^\nu_{\alpha I})^*$, resulting in a non-diagonal, complex, symmetric Dirac-Majorana mass matrix for neutrinos [46]:

$$\mathcal{L}_{\mathrm{DM}} = -\frac{1}{2}(\nu_{\mathrm{L}}^T\ \nu_{\mathrm{R}}^T)\begin{pmatrix} 0 & m_D^T \\ m_D & M_R \end{pmatrix}\begin{pmatrix} \nu_{\mathrm{L}} \\ \nu_{\mathrm{R}} \end{pmatrix} + \mathrm{h.c.} \tag{1.26}$$

where $M_R = \mathrm{diag}\,(M_I \dots)$ and the various flavors have been combined into vectors: $\nu_{\mathrm{L}} = (\nu_{\mathrm{L}\alpha} \dots)$, $\alpha = 1,2,3$ and $\nu_{\mathrm{R}} = (\nu_{\mathrm{R}I} \dots)$, $I = 1 \dots \mathcal{N}$. Using the Takagi factorization [47], the mass matrix can be brought to a diagonal form:

$$V^{\nu T} M_{\mathrm{DM}} V^\nu = \mathrm{diag}\,(m_i \dots) \qquad i = 1, \dots, 3 + \mathcal{N}. \tag{1.27}$$

Assuming that the charged lepton Yukawa couplings $Y^e_{\alpha\beta}$ are initially diagonal, the required unitary transformation of the neutrino fields is:

$$\begin{pmatrix} \nu_{\mathrm{L}} \\ \nu_{\mathrm{R}} \end{pmatrix} = V^\nu n \tag{1.28}$$

where $n = (n_1 \dots n_{3+\mathcal{N}})$ denotes the mass eigenstates, resulting in the Majorana mass term:

$$\mathcal{L}_{\mathrm{DM}} = -\frac{m_i}{2}(n_i n_i + n_i^\dagger n_i^\dagger) \qquad i = 1, \dots, 3 + \mathcal{N}. \tag{1.29}$$

[12] This section partially overlaps with (and improves upon) section 2.2.1. It is reproduced here for completeness, and the notations have been made consistent with the rest of chapter 1.

[13] One can always diagonalize the Majorana mass matrix with a suitable unitary transformation of the $\nu_{\mathrm{R}I}$ fields. Here we assume it to be in diagonal form already.

In the limit $|M_R| \gg |m_D|$, we can use an approximate block factorization, leading to the mass sub-matrices:

$$m_{\alpha\beta} \cong - \sum_{I=1\ldots\mathcal{N}} \frac{(m_D)_{\alpha I}(m_D)_{\beta I}}{M_I} \tag{1.30}$$

$$m_{IJ} \cong M_I \delta_{IJ} \tag{1.31}$$

and to the mass eigenstates $\nu = (n_1, n_2, n_3)$ and $N = (n_4 \ldots n_{3+\mathcal{N}})$ mixing with the flavor fields $\nu_{\mathrm{L}\alpha}$ and $\nu_{\mathrm{R}I}$ as follows:

$$\nu_{\mathrm{R}I} \cong N_I \tag{1.32}$$

$$\nu_{\mathrm{L}\alpha} \cong V^{\mathrm{PMNS}}_{\alpha i} \nu_i + \Theta_{\alpha I} N_I \tag{1.33}$$

$$\Theta_{\alpha I} \cong M_I^{-1} (m_D)_{\alpha I} \tag{1.34}$$

$V^{\mathrm{PMNS}}_{\alpha i}$ is numerically equal to the usual PMNS matrix, and the parameters $\Theta_{\alpha I}$ (called *mixing angles*) represent the mixing between the HNLs N_I and the "active" flavor states $\nu_{\mathrm{L}\alpha}$, which are charged under $SU(2)_{\mathrm{L}}$ (by opposition to the "sterile" flavor states $\nu_{\mathrm{R}I}$ which are singlets). It is due to this mixing that HNLs interact like heavy neutrinos with suppressed couplings.

If we take the Yukawa couplings to be uncorrelated and $\mathcal{O}(1)$, as originally assumed, then eq. (1.30) leads to light neutrino masses of order $v^2/|M_R|$. For M_R about some GUT scale $\sim 10^{15}\,\mathrm{GeV}$, this results in $|m_{\alpha\beta}| \lesssim 10^{-1}\,\mathrm{eV}$, in agreement with current bounds [26]. In the past decade, a class of *low-scale* seesaw models have risen in popularity, such as the νMSM [35], not least because of their falsifiability at existing or proposed experiments. In these models, M_R is postulated to be around the electroweak scale. The smallness of the light neutrino masses is then achieved either through small Yukawa couplings $|Y^\nu_{\alpha I}| \sim \frac{1}{v}\sqrt{|m_{\alpha\beta}||M_R|}$, which are too small to be accessible at current experiments, or through a cancellation between the terms of the sum in eq. (1.30). Such a cancellation (which could e.g. arise from a new symmetry) can be implemented by having two HNLs form a quasi-Dirac pair [48], i.e. be almost degenerate in mass, with mixing angles related by $\Theta_{\alpha 2} \approx \pm i \Theta_{\alpha 1}$. This would result in an *approximate* conservation of the total lepton number, the consequences of which will be reviewed in sections 1.2.4 and 1.3.

1.2.3 Parameter space of HNLs

The type-I seesaw Lagrangian (1.25) introduces a significant number of new parameters into the model: a Majorana mass M_I for each HNL, and a $3 \times \mathcal{N}$ complex matrix $Y^\nu_{\alpha I}$ of Yukawa couplings. In the case of the νMSM, this translates into six complex and two real parameters which are relevant to laboratory searches: the mixing matrix elements $\Theta_{\alpha I}$ and the masses M_2 and M_3 of the two heaviest HNLs. It is convenient to redefine the latter in terms of the central mass M and the physical mass splitting δM, such that $M_2 = M - \frac{\delta M}{2}$ and $M_3 = M + \frac{\delta M}{2}$. In this model, the mixing angles associated with the lightest HNL must be tiny[14], causing it to decouple. Due to eq. (1.34), this requires

[14] In the νMSM, the lightest HNL, N_1, is a dark matter candidate. The magnitude of its Yukawa couplings is thus strongly constrained by Lyman-α and X-ray observations [49, 50]. This indirectly sets an upper bound on the mass of the lightest neutrino.

one of the light neutrinos to be much lighter than the other two. All physical masses are thus constrained.

Not all these parameters are independent. Not only they are related through the Takagi factorization (1.27), but they must also reproduce the observed light neutrino mass-squared differences and the PMNS matrix (i.e. the PMNS matrix must be embedded in the extended unitary transformation matrix V^ν which appears in the Takagi factorization). In particular, for large HNL mixing angles — much larger than the naive seesaw expectation of $|\Theta|^2 \sim |m_{\alpha\beta}|/|M_\mathrm{R}|$ — eq. (1.30), combined with eq. (1.34), leads to particularly strong constraints. In order to keep the light neutrino masses $m_{\alpha\beta}$ light, a cancellation must take place in the right-hand side. For two HNLs $N_{2,3}$ and in the limit of massless neutrinos ($m_{\alpha\beta} = 0$), this cancellation must be exact, and it has been shown in ref. [48] that the only solution that is stable under radiative corrections is:[15]

$$M_3 = M_2 \quad \text{and} \quad \Theta_{\alpha 3} = \pm i \Theta_{\alpha 2} \text{ for all } \alpha. \tag{1.35}$$

As will be discussed in section 1.2.4, in this limit the lepton number violating effects coming from the HNL Majorana masses cancel out exactly, leading to a phenomenology similar to that of a Dirac fermion. For this reason, the two HNLs are said to form a *quasi-Dirac* HNL pair.

Going back to the general case, a convenient method to express all physical, non-equivalent choices of parameters was provided in ref. [51] (see also ref. [52] for a detailed derivation). This Casas-Ibarra parametrization relates the HNL mixing angles $\Theta_{\alpha I}$ to the light neutrino masses $\hat{M}_\nu = \mathrm{diag}\,(m_1, m_2, m_3)$, the HNL masses $\hat{M}_N = \mathrm{diag}(M - \frac{\delta M}{2}, M + \frac{\delta M}{2})$ and the PMNS matrix V_PMNS (including one Majorana phase[16] η) through the matrix equation:[17]

$$\Theta = i V_\mathrm{PMNS} \hat{M}_\nu^{\frac{1}{2}} \Omega \hat{M}_N^{-\frac{1}{2}} \tag{1.36}$$

where Ω is a 3×2 matrix in which is embedded an arbitrary 2×2 complex orthogonal matrix, such that $\Omega^T \Omega = \mathbb{1}_{2\times 2}$. Depending on the neutrino hierarchy (normal: NH or inverted: IH), such a matrix can be parametrized as:

$$\Omega = \begin{pmatrix} 0 & 0 \\ \cos(\omega) & \sin(\omega) \\ -\xi \sin(\omega) & \xi \cos(\omega) \end{pmatrix} \quad (\mathrm{NH}) \quad \text{or} \quad \begin{pmatrix} \cos(\omega) & \sin(\omega) \\ -\xi \sin(\omega) & \xi \cos(\omega) \\ 0 & 0 \end{pmatrix} \quad (\mathrm{IH}) \tag{1.37}$$

where ω is a complex angle, and $\xi \in \pm 1$ is a parity which can be set to $+1$ without losing generality[18]. We are thus left with one complex free parameter, ω, and three real ones: the HNLs mass M, their mass splitting δM, and the Majorana phase η. The remaining parameters are not specific to HNLs, but rather represent input from neutrino oscillations; as such, they can be measured

[15] Keeping the νMSM notations, with indices $I = 2, 3$ denoting the two heaviest HNLs.

[16] The PMNS matrix for Majorana neutrinos is related to the usual one by redefining $V_\mathrm{PMNS} \to V_\mathrm{PMNS} \cdot \mathrm{diag}\,(1, e^{i\eta}, 1)$, where η is the Majorana phase (there is only one such phase when considering two HNLs). For $\mathcal{N} > 2$ HNLs, additional Majorana phases appear.

[17] This equation is formally valid for all $\mathcal{N}$, for suitable definitions of $\hat{M}_N$ and Ω.

[18] The case $\xi = -1$ can then be recovered by simultaneously redefining $\omega \to -\omega$ and $N_3 \to -N_3$ [53].

or constrained. We estimate them using the global fit from the latest NuFIT publication [54, 55].

Armed with the Casas-Ibarra parametrization, we can then proceed to scan the parameter space allowed by neutrino oscillations. For convenience, we often define $X_\omega = e^{\mathrm{Im}(\omega)}$. This parameter is closely related to the magnitude of the HNL mixing angles, which aro minimal for $\mathrm{Im}(\omega) = 0$, and behave asymptotically as $|\Theta_{\alpha I}| \sim e^{|\mathrm{Im}(\omega)|}$ for large $|\mathrm{Im}(\omega)|$ [37]. It is noteworthy that we automatically recover the quasi-Dirac limit when $M_2 \approx M_3$ and $|\,\mathrm{Im}(\omega)| \gg 1$. In this limit, and considering for example the νMSM with two HNLs, the combinations of HNL mixing angles which are compatible with neutrino oscillation data are quite constrained. This is represented in figure 1.1 using a ternary plot, which shows all the possible ratios of $|\Theta_{eI}|^2 : |\Theta_{\mu I}|^2 : |\Theta_{\tau I}|^2$ ($I = 2$ or 3 since the magnitudes are the same) as a function of the neutrino hierarchy and the level of agreement with the NuFIT 5.0 global fit.[19] This plot was obtained by scanning over the neutrino oscillation parameters within a $\Delta\chi^2$ (with respect to the best-fit point for the given hierarchy) corresponding to the specified confidence level, and then scanning over the remaining free parameters of the Casas-Ibarra parametrization. It is remarkable that the ratios corresponding to “single-flavor mixing” — often used by experiments to report their sensitivity or limits — which are represented by the three vertices of the triangle, are incompatible with neutrino oscillation data within this model.

1.2.4 Lepton number violation

Due to their mixing with neutrino flavor fields, HNLs are unstable particles. Here we shall focus on the case where they are weakly produced on their mass shell, typically in the decay of a heavy meson or vector boson, and then decay weakly. The Feynman rules for Majorana HNLs contain four different propagators [2]. They are represented in figure 1.2 for a generic electroweak process, along with the mixing angles corresponding to both interaction vertices. Two of them (figures 1.2a and 1.2b) conserve the total lepton number, like Dirac fermions, and lead to *lepton number conserving* (**LNC**) processes, while two others (figures 1.2c and 1.2d) violate it by two units — leading to *lepton number violating* (**LNV**) processes — and are only possible if HNLs are Majorana fermions. The flow of lepton number, as well as the various charges, can be inferred from the direction of the arrows.

Suppose now that $\mathcal{N}$ HNLs are nearly degenerate in mass M and satisfy eq. (1.30). Additional HNLs can be omitted if they are sufficiently decoupled, like e.g. N_1 in the νMSM. All HNLs taking part in the process are therefore almost indistinguishable, except for their small mass splittings. If those are small enough (such that coherence is maintained by the production, propagation and decay processes [22]), the transition amplitudes for different propagating HNLs will interfere, and need to be summed when computing the total amplitude for the decay chain. If we were to omit the space-time-dependent part of the amplitude introduced by the phase $e^{-iq_I\cdot x}$ in the HNL propagator, then, factoring out the common part $\hat{A}^{\pm\pm}_{\alpha\beta}$ from the amplitude[20] (since it is nearly equal for all

[19] Without the Super-Kamiokande atmospheric data.

[20] The kinematical variables have been omitted for brevity.

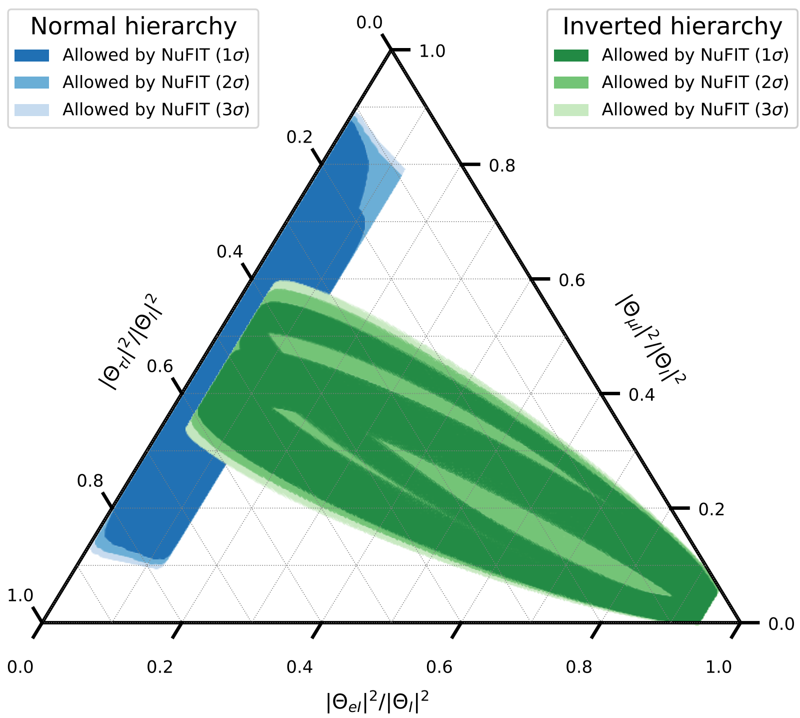

Figure 1.1: Representation of the set of HNL mixing angles $\Theta_{\alpha I}$ which are compatible with the neutrino oscillation data [54, 55] at various levels, within the νMSM with two HNLs forming a quasi-Dirac pair. $|\Theta_I|^2$ is defined as $\sum_{\alpha=e,\mu,\tau} |\Theta_{\alpha I}|^2$ and $I = 2$ or 3.

the nearly degenerate HNLs), we would obtain for the total amplitude:

$$\mathcal{A}^{\pm\pm}_{\alpha\beta,\mathrm{tot}} = \left(\sum_I \Theta^{\pm}_{\alpha I} \Theta^{\pm}_{\beta I} \right) \hat{A}^{\pm\pm}_{\alpha\beta} \tag{1.38}$$

where $-$ denotes the HNL interacting through its left chiral component (N_I), $+$ through its right one ($N_I^\dagger$), α and β are the flavors through which it respectively interacts during production and decay [21], and we have used the compact notation $\Theta^-_{\alpha I} = \Theta_{\alpha I}$ and $\Theta^+_{\alpha I} = \Theta^*_{\alpha I}$. The total amplitude for a LNV process is therefore proportional to $\sum_I \Theta_{\alpha I} \Theta_{\beta I}$ or its complex conjugate. However, substituting eq. (1.34) into eq. (1.30) and factoring out the common mass leads to:

$$m_{\alpha\beta} \cong -M \sum_I \Theta_{\alpha I} \Theta_{\beta I} \tag{1.39}$$

The empirical requirement that $|m_{\alpha\beta}| \lll M$ thus implies that the mixing angles must satisfy $\left| \sum_I \Theta_{\alpha I} \Theta_{\beta I} \right| \lll 1$, a priori leading us to the conclusion that

[21] For instance, for a LNC process involving the sub-diagram in figure 1.2a, where the HNL mixes with flavor α through it right chiral component (+) during production, and with flavor β through its left chiral component ($-$) during decay, the total amplitude will be denoted by $\mathcal{A}^{+-}_{\alpha\beta,\mathrm{tot}}$.

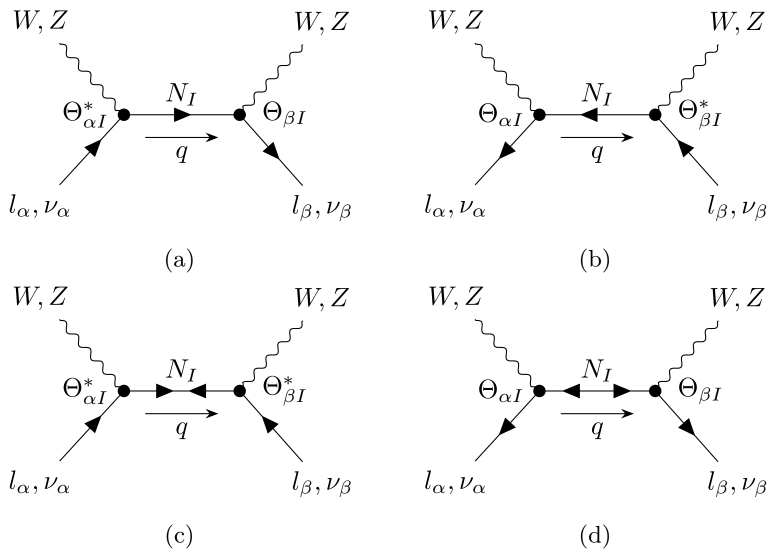

Figure 1.2: The four propagators involved in the production in $\bar{x}_P$ (left vertex), propagation and then destruction in $\bar{x}_D$ (right vertex) of a Majorana HNL. The charge/helicity of the accompanying leptons can be inferred from the direction of the arrows.

LNV amplitudes must be suppressed in order for light neutrinos to remain light, as found in [48] in a more general case. LNC amplitudes being proportional to $\sum_I \Theta^*_{\alpha I}\Theta_{\beta I}$ or its complex conjugate, this argument does not apply to them. In the next section, we shall see how including the space-time-dependent phase in the amplitude affects our conclusion.

1.3 Coherent HNL oscillations

If we reintroduce the space-time-dependent phase $e^{-iq_I \cdot x}$ into the amplitude, we immediately see that multiple mass eigenstates will produce interference terms of the form $e^{i(q_J - q_I)\cdot x}$, resulting in space-time-dependent interference, i.e. oscillations. However, the precise form taken by these oscillations is not initially obvious. Although HNL oscillations share a similar origin with neutrino flavor oscillations, we can reasonably expect their phenomenology to be quite different, since HNLs are heavy, unstable particles coming from the addition of sterile / right-handed flavor states. In particular, while lepton number violating effects are kinematically suppressed in the interactions of light neutrinos (if they are Majorana particles), there is no such suppression in decays of on-shell particles [56, 57], and we therefore expect some HNL decays to be lepton number violating. We also need to ensure that no decoherence takes place, as it would destroy the oscillations. In order to determine the precise form of these oscillations, it seems reasonable to follow an ab-initio approach based on quantum field theory and on existing knowledge about oscillations of unstable particles. After a brief (and non-exhaustive) review of the existing literature in section 1.3.1, the external wave packet model is employed in section 1.3.2 to derive a description of HNL oscillations. Sections 1.3.3 to 1.3.5 are then dedicated to the phenomenology of these oscillations, and discuss some interesting limits and implications.

1.3.1 Literature review

Neutrino-antineutrino oscillations have been considered ever since Pontecorvo's seminal article [17], by analogy with meson-antimeson oscillations. It was only later that the possibility of flavor oscillations was contemplated [18]. In these first articles, no explicit model was ever considered, but the implicit assumption was that neutrinos were Majorana particles [58]. In ref. [59], the standard Dirac-Majorana mass term made an apparition, and a formula for ν_e disappearance was derived, and then generalized in ref. [60] to obtain the standard formula of neutrino oscillations (minus CP violation). The focus had then moved to flavor oscillations, and neutrinos were still assumed to be ultra-relativistic. The introduction of the seesaw mechanism [27–34] popularized the idea of heavy Majorana neutrinos. Unfortunately, in its initial form, it resulted in GUT-scale neutrinos which were way out of the experimentally testable realm. Although oscillations of heavy, unstable particles have been studied in the literature [22], the aim was then mostly to describe heavy meson oscillations, and no phenomenological study of HNL oscillations was done until very recently. The growing popularity in the past decade of the νMSM as a viable BSM theory has put sub-electroweak-scale HNLs under the spotlight, and led to a renewed interest in HNL oscillations. A first study of flavor oscillations of HNLs was conducted in ref. [61], but helicity was averaged out and the formalism was not very suitable for calculations using S-matrix scattering theory. A similar study was conducted in ref. [62] but, while technically correct, the formalism is not straightforward to reconcile with the S-matrix one either, and lepton-number-violating (LNV) oscillations were only considered in the specific case where HNLs form a quasi-Dirac pair (CP eigenstate). To the best of the author's knowledge, the first *phenomenological* studies of HNL oscillations are refs. [63] and [64]. The first study of HNL oscillations in relation to LHC experiments is ref. [65], which focused again on the quasi-Dirac limit (hence neglecting CP violation), and derived an observable R_{ll} based on the same-sign to opposite-sign dilepton ratio, which can take non-trivial values ($R_{ll} \neq 0, 1$) when the HNL lifetime is of the same order of magnitude as the oscillation frequency of the quasi-Dirac pair. The study was limited to the case of the inverse seesaw in a left-right symmetric model. Ref. [66] additionally considered the linear seesaw, and computed regions where R_{ll} is non-trivial for specific models. Resolvable HNL oscillations were also discussed briefly. A more comprehensive study of resolvable HNL oscillations, in the context of the νMSM, can be found in ref. [67]. Finally, in ref. [68], the angular distribution of the HNL decay products was studied in an e^+e^- collider setting, as a function of the mass splitting δm of the quasi-Dirac pair. As we shall see in chapter 2, this effect can be understood as the kinematical counterpart to the R_{ll} ratio. Further phenomenological studies of HNL oscillations can be found in refs. [56, 69–72].

The observation of HNL oscillations would prove the existence of not only one, but (at least) two nearly-degenerate HNLs, which would be a smoking gun for the ARS mechanism. By allowing us to measure the HNL mixing angles and relevant phases, it would shed light on the structure of the extended neutrino sector and provide a stringent test of the seesaw mechanism. Moreover, HNL

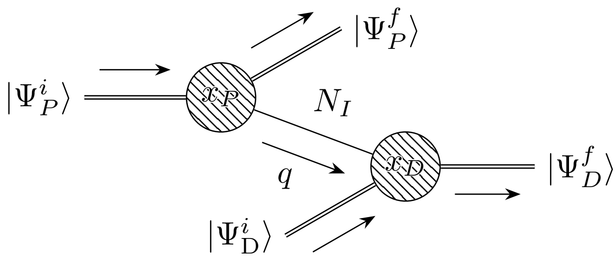

Figure 1.3: Schematic process with the intermediate HNLs and the external wave packets represented.

oscillations must be allowed for when carrying out displaced vertex searches at colliders. A complete, accurate and usable model of HNL oscillations, valid even near the seesaw bound (where CP-violating effects cannot be neglected), is therefore essential when studying HNLs.

1.3.2 The external wave packet model

In order to accurately describe processes involving multiple intermediate HNLs, it is essential to treat them *coherently*, i.e. compute a single amplitude for the entire process, comprising the HNL production, propagation and decay. Indeed, only external particles result in a projection of the quantum state of the system when they eventually interact with the environment. Since multiple intermediate HNLs have a different dispersion relation $q^2 = M_I^2$, they each contribute a different complex phase $e^{-iq\cdot x}$. When the phase shift between two mass eigenstates becomes sufficiently large, it can significantly alter the total amplitude, ultimately leading to oscillations. It is therefore fundamental to carefully treat the space-time dependence of the amplitude. In many models, HNLs are sufficiently long-lived to leave a displaced vertex when decaying. However, the usual S-matrix scattering theory is only suitable to describe plane waves, which are inherently delocalized. A natural framework to study HNL oscillations is thus the *external wave packet* model [73, 74] (see also [22, 75–77] and references therein for recent reviews), in which only external particles are treated as asymptotic states, and are localized by associating them with finite wave packets. Wave packets are required in order to obtain interference between multiple mass eigenstates while preserving exact momentum conservation at every vertex [76, 78]. We do not lose any generality by employing this framework, since it is a superset of the standard scattering theory. In this section, we shall briefly sketch the derivation of HNL oscillations within the external wave packet model, without carrying out the explicit computation. The interested reader is referred to chapter 4 from [22] for a detailed proof (see also the author's master thesis [67] or a recent, independent derivation in ref. [79]).

Consider now the generic process depicted in figure 1.3. Incoming particles are respectively denoted by $|\Psi_P^i\rangle$ and $|\Psi_D^i\rangle$ at the HNL production and destruction vertices x_P and x_D, and outgoing particles by $|\Psi_P^f\rangle$ and $|\Psi_D^f\rangle$. For HNL decays, $|\Psi_D^i\rangle = |0\rangle$. External particle states are assumed to be uncorrelated, such that the initial and final states $|\Psi^i\rangle$ and $|\Psi^f\rangle$ can be written as tensor products of the individual particle states $|\Psi_1\rangle \otimes |\Psi_2\rangle \otimes \ldots$ (anti-/symmetrized as needed). Blobs denote the processes involved in the HNL production or destruction, which only trivially depends on N_I, i.e. through the mixing angles $\Theta_{\alpha I}$ and kinematic variables. Each external particle state $|\Psi_j\rangle$ can be decomposed into Fourier modes $|\mathbf{k}_j\rangle$:

$$|\Psi_j\rangle = \int \mathrm{d}\Omega_{\mathbf{k}_j} \psi_j(\mathbf{k}_j) \, |\mathbf{k}_j\rangle \tag{1.40}$$

where $\mathrm{d}\Omega_{\mathbf{k}_j}$ is the Lorentz-invariant integration measure and $\psi_j(\mathbf{k}_j)$ is the *wave packet* of the particle (i.e. its one-particle wave-function in the asymptotic, non-interacting limit). This allows expressing the total transition amplitude $\mathcal{A}(\Psi^i \to \Psi^f)$ in terms of plane-wave amplitudes $\mathcal{A}_I^{\mathrm{PW}}(\{\mathbf{k}_j\})$, which can be readily computed with the usual methods:

$$\begin{aligned}\mathcal{A}(\Psi^i \to \Psi^f) &= \langle \Psi^f | \hat{S} | \Psi^i \rangle \\ &= \int \mathrm{d}\Omega_{\mathbf{k}_1^f \dots \mathbf{k}_m^f} \mathrm{d}\Omega_{\mathbf{k}_1^i \dots \mathbf{k}_n^i} \psi_1^{f*}(\mathbf{k}_1^f) \dots \psi_1^i(\mathbf{k}_1^i) \dots \cdot \sum_I \mathcal{A}_I^{\mathrm{PW}}(\mathbf{k}_1^i \dots \mathbf{k}_n^i \to \mathbf{k}_1^f \dots \mathbf{k}_m^f)\end{aligned} \tag{1.41}$$

For brevity we have omitted the dependence on external spins, which can be factored out for nearly-degenerate HNLs (cf. ref. [22], sec. 4.3), but we have kept the sum over mass eigenstates N_I. We have also included the mixing angles $\Theta_{\alpha I}$ in $\mathcal{A}_I^{\mathrm{PW}}$. Equation (1.41) does not explicitly depend on any space-time coordinates. We can reveal this dependence by noticing that plane waves from eq. (1.40) only lead to constructive interference in a small region of space-time near the trajectory of each particle. The product of all wave packets is thus non-vanishing only in the vicinity of x_P or x_D, where they overlap. Following [22], we can without losing generality work with zero-centered wave packets $\bar{\psi}_j(k_j) = e^{-ik_j \cdot x_{D/P}} \psi_j(k_j)$, which simplify the definition of the overlap region. Shifting the wave packets introduces a complex phase $e^{\mp ik_j \cdot x_P}$ for particles interacting at x_P and $e^{\mp ik_j \cdot x_D}$ for those interacting at x_D, where the $-$ sign is for incoming particles and the $+$ sign for outgoing ones. Using momentum conservation, this results in an overall phase factor of $e^{-iq\cdot(x_D - x_P)}$.

Let us now focus on the case where the separation between the production and decay vertex is much larger than the wave packet width. This is always the case if the HNL leaves a displaced vertex, but it might still be true even if the vertex is not experimentally resolvable. As we shall see in section 1.3.5, this implies that the coherence production / detection condition is satisfied. In this case, the residual space-time dependence, of the order of the width σ_x of the overlap region, can be shown to be negligible. For sufficiently degenerate HNLs, we can factor out the kinematical part of the amplitude from the sum over I, leaving only the HNL propagators, the mixing angles $\Theta_{\alpha I}$ and the phase $e^{-iq\cdot(x_D - x_P)}$. We can finally perform the integration over external momenta $\mathbf{k}_j$. For on-shell HNLs, the only contribution to the integral comes from the pole in each propagator, which imposes the dispersion relation $q^2 = M_I^2 - iM_I\Gamma_I$. Combined with momentum conservation at each vertex, it effectively results in an integration over the 3-momentum $\mathbf{q}$, where $E_I(\mathbf{q})$ is treated as a dependent variable. The different dispersion relations lead to a systematic phase shift $e^{-i(q_I - q_J)\cdot(x_D - x_P)}$ between the contributions of two HNLs N_I and N_J to the integrand, where we have defined $q_I \equiv (E_I(\mathbf{q}), \mathbf{q})$. Since HNLs are heavy, the phase (including the absorptive part) can in good approximation be evaluated in their common rest frame, yielding $e^{-i(M_I-M_J)\tau-\frac{\Gamma_I-\Gamma_J}{2}\tau}$, where $\tau=\sqrt{(x_D-x_P)^2}$ is the proper time between the HNL production and decay vertex. If wave packets are sufficiently peaked in momentum space (but not so much as to allow us to determine the mass eigenstate), we can finally factor out the slowly-varying part of the phase from the integral. Using the same compact notation as in eq. (1.38):

$$\mathcal{A}^{\pm\pm}_{\alpha\beta}(\Psi^i\to\Psi^f)=\left(\sum_{I=1}^{\mathcal{N}}\Theta^{\pm}_{\alpha I}\Theta^{\pm}_{\beta I}e^{-i(M_I-\hat{M})\tau-\frac{\Gamma_I}{2}\tau}\right)\cdot\hat{A}^{\pm\pm}_{\alpha\beta}(\Psi^i\to\Psi^f) \quad (1.42)$$

where $\hat{A}^{\pm\pm}_{\alpha\beta}(\Psi^i\to\Psi^f)$ is the common part of the amplitude, for a single HNL of reference mass $\hat{M}\approx M_I$, and without the absorptive part nor the mixing angles. In practical calculations, this one-HNL amplitude can easily be evaluated using the usual plane-wave scattering theory. Taking the absolute square of the amplitude and introducing the relevant phase-space factors, we then obtain the relation between the coherent differential event rates (or cross-sections) for $\mathcal{N}$ HNLs and 1 HNLs:

$$\mathrm{d}\Gamma^{\pm\pm}_{\alpha\beta}(\tau)=\left|\sum_{I=1}^{\mathcal{N}}\Theta^{\pm}_{\alpha I}\Theta^{\pm}_{\beta I}e^{-iM_I\tau-\frac{\Gamma_I}{2}\tau}\right|^2\mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta} \quad (1.43)$$

where $\mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta}$ is computed for a single HNL with unit mixing angles and no absorptive part $e^{-\Gamma\tau}$. Let us emphasize that $\mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta}$ may in general result in very different kinematics for the various combinations of chiralities $\pm$. If we were to omit the interference terms (for instance if coherence is lost), we would obtain instead the incoherent width:

$$\mathrm{d}\Gamma^{\pm\pm}_{\alpha\beta,\mathrm{inc}}(\tau)=\left(\sum_{I=1}^{\mathcal{N}}\left|\Theta_{\alpha I}\right|^2\left|\Theta_{\beta I}\right|^2e^{-\Gamma_I\tau}\right)\mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta} \quad (1.44)$$

Finally, let us briefly mention the case where the wave packets are much wider (in position space) than the separation between the production and decay vertex. This corresponds to a prompt HNL decay, and wave packets can be well-approximated as plane waves, up to negligible boundary effects. In this limit, there are no oscillations, but the finite decay width Γ_I (which is encoded in the self-energy of the propagator) will typically introduce a dependence of the final amplitude on ${}^{\delta M}/_{\Gamma}$ and ${}^{\delta\Gamma}/_{\Gamma}$. This case is already well-covered in the literature; see for instance [68,80]. Interestingly, the plane-wave width matches the τ-integrated coherent width, i.e. integrating over space-time at the amplitude or squared amplitude level yields the same result. This suggests a smooth transition between the two regimes when looking only at space-time integrated observables. Incidentally, it also explains why [68] obtains the correct result while using the plane-wave formalism despite considering displaced HNLs.

1.3.3 Oscillations of quasi-Dirac HNLs

Let us now specialize eq. (1.43) to the case of two HNLs N_1 and N_2 forming a quasi-Dirac pair (introduced in section 1.2.3), i.e. $M_1=M-\frac{\delta M}{2}$, $M_2=M+\frac{\delta M}{2}$

and $\Theta_{\alpha 2} = \pm i \Theta_{\alpha 1}$. In this limit, eq. (1.43) becomes:

$$\mathrm{d}\Gamma^{\pm\pm}_{\alpha\beta}(\tau) = 2\left(1 \pm \cos\left(\delta M \tau\right)\right) e^{-\Gamma\tau} \mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta} \tag{1.45}$$

where the $+$ sign is for lepton-number-conserving processes ($\mathrm{d}\Gamma^{+-}_{\alpha\beta}$ and $\mathrm{d}\Gamma^{-+}_{\alpha\beta}$), and the $-$ sign for lepton-number-violating ones ($\mathrm{d}\Gamma^{++}_{\alpha\beta}$ and $\mathrm{d}\Gamma^{--}_{\alpha\beta}$). The corresponding τ-integrated differential width is:

$$\int_0^\infty \mathrm{d}\tau \, \mathrm{d}\Gamma^{\pm\pm}_{\alpha\beta}(\tau) = \left(\frac{1}{\Gamma} \pm \frac{\Gamma}{\Gamma^2 + (\delta M)^2}\right) \mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta} \tag{1.46}$$

Integrating over phase-space and considering the "standard" LNV / LNC ratio R_{ll}, we recover the usual expression:

$$R_{ll} = \frac{(\delta M)^2}{2\Gamma^2 + (\delta M)^2} \tag{1.47}$$

In order to clarify when lepton-number violation is possible, let us now consider various limits. If the HNL is short-lived (i.e. it is observed at $\Gamma\tau \gtrsim 1$), only the integrated rate can be measured:

1. If $\Gamma \gg \delta M$, then the HNL pair behaves like a single Dirac HNL[22] and $R_{ll} \to 0$ (the HNLs do not have time to oscillate before decaying).
2. If $\Gamma \ll \delta M$, then the HNL pair behaves like a single Majorana HNL[23] and $R_{ll} \to 1$ (the HNLs undergo many oscillations before decaying).
3. For $\Gamma \sim \delta M$, non-trivial R_{ll} ratios are possible.

If the HNL is long-lived (i.e. it is observed at $\Gamma\tau \ll 1$), then its behaviour depends not only on the mass splitting δM, which is an intrinsic parameter of the theory, but also on the range of proper time τ accessible at a given experiment:

1. If $\delta M \tau \ll 2\pi$, the HNL pair is observed before the onset of oscillations, and it behaves like a single Dirac HNL, with $R_{ll} \to 0$.
2. If $\delta M \tau \gg 2\pi$, oscillations are averaged out, and the HNL pair behaves like a single Majorana HNL, with $R_{ll} \to 1$.
3. If $\delta M \tau \sim 2\pi$, then oscillations must be accounted for. If it is possible to experimentally reconstruct, for each event, the proper time $\tau = \sqrt{(x_D - x_P)^2}$ between the production and decay vertex of the HNL, then oscillations can be resolved, i.e. the τ-differential event rate will show a periodic modulation according to eq. (1.45). These oscillations can be interpreted as particle-antiparticle oscillations between a Dirac fermion N and its antiparticle $\bar{N}$ (like neutral meson oscillation).

We will from now on refer to case 1. as the **Dirac-like** limit of quasi-Dirac HNLs, and to case 2. as their **Majorana-like** limit, regardless of whether the HNLs are short-lived or long-lived.

These results are well known in the quasi-Dirac limit (see e.g. [66]), but eq. (1.43) makes it straightforward to generalize them to arbitrary numbers of

[22] With a $4\times$ enhancement of all cross-sections and a lifetime only half as long.
[23] With a $2\times$ enhancement of all cross-sections and the same lifetime.

nearly-degenerate HNLs and arbitrary mixing angles. It is also interesting to notice that HNLs may behave as Dirac in one experiment and as Majorana in another, if the two experiments probe sufficiently different τ ranges. This could in principle allow to bisect δM without explicitly observing oscillations.

1.3.4 CP violation

Recently, there has been a strong interest in observing CP-violating HNL oscillations [69–72, 80, 81]. In this section, we aim to clarify whether CP violation is possible at all for quasi-Dirac HNLs, if it is allowed in LNC channels, and which models can give rise to it. To this end, let's introduce an additional complex phase between the mixing angles of the two HNLs: $\Theta_{\alpha 2} = \pm i e^{i\theta_\alpha} \Theta_{\alpha 1}$. To remain generic, we let it depend on flavor. The coherent width becomes (where for brevity we have included the exponential decay in the one-HNL width $\mathrm{d}\tilde{\Gamma}^{\pm\pm}_{\alpha\beta}(\tau) = \mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta} e^{-\Gamma\tau}$):

$$\begin{aligned}
\mathrm{d}\Gamma^{+-}_{\alpha\beta}(\tau) &= 2\left(1 + \cos\left(\delta M\tau + \theta_\alpha - \theta_\beta\right)\right) \mathrm{d}\tilde{\Gamma}^{+-}_{\alpha\beta}(\tau) \\
\mathrm{d}\Gamma^{-+}_{\alpha\beta}(\tau) &= 2\left(1 + \cos\left(\delta M\tau - \theta_\alpha + \theta_\beta\right)\right) \mathrm{d}\tilde{\Gamma}^{-+}_{\alpha\beta}(\tau) \\
\mathrm{d}\Gamma^{++}_{\alpha\beta}(\tau) &= 2\left(1 - \cos\left(\delta M\tau + \theta_\alpha + \theta_\beta\right)\right) \mathrm{d}\tilde{\Gamma}^{++}_{\alpha\beta}(\tau) \\
\mathrm{d}\Gamma^{--}_{\alpha\beta}(\tau) &= 2\left(1 - \cos\left(\delta M\tau - \theta_\alpha - \theta_\beta\right)\right) \mathrm{d}\tilde{\Gamma}^{--}_{\alpha\beta}(\tau)
\end{aligned}$$

We immediately see that the phases θ_α are CP-violating. If the phase $\theta_\alpha \equiv \theta_{\mathrm{LV}}$ does not depend on the generation, then it only leads to CP violation in LNV channels. But this assumption does not hold in general. We may wonder if it is possible to keep the light neutrinos light when introducing CP violation. This can be studied directly by making use of eq. (1.39).

$$m_{\alpha\beta} \cong -M\Theta_{\alpha 1}\Theta_{\beta 1}\left(1 - e^{i(\theta_\alpha + \theta_\beta)}\right) \tag{1.48}$$

Consider the Frobenius norm $\mathrm{tr}\left(m_\nu^\dagger m_\nu\right)$ of the light neutrino mass matrix. In the mass basis:

$$\mathrm{tr}\left(m_\nu^\dagger m_\nu\right) = \sum_{i=1}^{3} m_i^2 \tag{1.49}$$

while in the flavor basis:

$$\mathrm{tr}\left(m_\nu^\dagger m_\nu\right) = \sum_{\alpha,\beta} m^*_{\alpha\beta} m_{\alpha\beta} = 2M^2 \sum_{\alpha,\beta} \left|\Theta_{\alpha 1}\right|^2 \left|\Theta_{\beta 1}\right|^2 \left(1 - \cos(\theta_\alpha + \theta_\beta)\right) \tag{1.50}$$

Each term in the sum is positive definite, therefore if any $\theta_\alpha + \theta_\beta \neq 0$, the norm, and therefore at least one light neutrino mass, will be non-zero. In order to keep the light neutrino masses small, the phases must satisfy $\theta_\alpha + \theta_\beta \cong 0$ for all α and β, and in particular for $\alpha = \beta$. The only solution is that all $\theta_\alpha \cong 0$, which directly shows that CP violation must be suppressed for quasi-Dirac HNLs. More precisely, the knowledge of the light neutrino masses allows us to place an upper bound on CP-violating phases:

$$|\theta_\alpha| \lesssim \frac{|m_{\alpha\beta}|}{2M\,|\Theta_{\alpha 1}|\,|\Theta_{\beta 1}|} \tag{1.51}$$

This expression tells us that no matter how large the mixing angles are, the absolute scale of CP violation will always remain approximately the same, around the naive seesaw expectation $|m_\nu|/M$. With only two HNLs, CP violation should therefore be experimentally unobservable. It has nonetheless been shown to still be sufficient for producing the baryon asymmetry of the Universe [37, 82].

The case of three quasi-degenerate HNLs turns out to be qualitatively different. Using the Casas-Ibarra parametrization described in section 1.2.3 (but for three HNLs), it is straightforward to show that "typical" choices of mixing angles produce non-zero CP-odd phases $\varphi^{(3)}_{\alpha\beta,IJ} = \arg(\Theta^{\pm}_{\alpha I}\Theta^{\pm}_{\beta I}\Theta^{\mp}_{\alpha J}\Theta^{\mp}_{\beta J})$ for all combinations of flavors α and β, and $I \neq J = 1, 2, 3$, resulting in CP violation in lepton number violating decays. Therefore, detecting CP violation in HNL oscillations or decays would indicate the presence of at least three right-handed neutrinos.

1.3.5 Coherence conditions

In order for coherent oscillations to be possible, two conditions must be satisfied, as discussed in ref. [22,77]. The coherent production / detection condition is the requirement that a precise measurement of the external momenta (at either vertex) should not allow to distinguish which particular mass eigenstate mediated a specific instance of the process. More technically, the energy and momentum spread of the wave packets must be sufficiently large so that the integration over external momenta hits all HNL poles in the propagator. As shown in ref. [83], due to the uncertainty principle, this condition has an intuitive counterpart in position space. It is equivalent to the statement that the width of the overlap region in position space is smaller than the oscillation wavelength. For all cases where oscillations (and not just their integrated effect) are of practical relevance, this condition is trivially satisfied due to the (classical) requirement that the vertexing resolution be smaller than the oscillation wavelength.

The second condition is the coherent propagation. It is the requirement that the wave packets do not separate due to their different group velocities, and as such it crucially depends on the temporal resolution of the detector [84]. In other words, it must not be possible to distinguish wave packets using a time-of-flight measurement. This condition is satisfied provided the observation distance L and energy width σ_E of the wave packets are such that:

$$2\pi \frac{L}{L_{\mathrm{osc}}} \ll \frac{E}{\sigma_E} \tag{1.52}$$

where L/L_{osc} is the number of oscillation periods. This condition is also trivially satisfied in all relevant cases, due to the (classical) requirement that the energy resolution $\delta E > \sigma_E$ must be sufficiently small in order to reconstruct the oscillating signal without smearing it. We thus see that the production, propagation and decay processes remain coherent in all situations where resolving oscillations would make any sense. Ambiguities may only arise for prompt HNLs, for which the above reasoning does not hold. However, in this case, we can only look at space-time-integrated observables which, as we noticed, do not depend on the coherence of the process.

1.4 Recap & introduction to the articles

After a short introduction to the Standard Model of particle physics in section 1.1.1, we have highlighted in section 1.1.2 a number of observational shortcomings in this otherwise very well tested model: non-zero neutrino masses, the matter-antimatter asymmetry, dark matter — and we have briefly reviewed the evidence supporting these observations.

Heavy neutral leptons were then introduced in section 1.2 as a minimal, yet powerful extension of the Standard Model, which could potentially resolve all of these problems simultaneously (as in the neutrino minimal standard model). After presenting the rationale behind this model in section 1.2.1 and introducing the type-I seesaw mechanism in section 1.2.2, sections 1.2.3, 1.2.4 and 1.3 discussed several aspects of the phenomenology of HNLs, such as their viable parameter space, the (non-)conservation of the total lepton number and the signatures resulting from multiple interfering HNLs (lepton number conservation and HNL oscillations). In particular, section 1.3 focused on HNL oscillations, which is the most distinctive feature of models featuring multiple nearly-degenerate HNLs. Using an ab-initio treatment based solely of QFT with wave packets, we obtained in section 1.3.2 an unambiguous description of HNL oscillations for any number of nearly-degenerate HNLs, with any mixing angles. The next sections then discussed several aspects of these oscillations, such as the simplified form that they take for quasi-Dirac HNLs (section 1.3.3), the presence or absence of CP violation (section 1.3.4) and the conditions under which the process remains coherent (section 1.3.5).

This introduction underscores how, despite being one of the simplest possible extensions of the Standard Model, heavy neutral leptons can nonetheless feature a very rich phenomenology, including coherent oscillations, CP violation and lepton number violation, all of which may be either present or suppressed depending on the specific parameters of the model. It is particularly noteworthy that all of these aspects of the phenomenology of HNLs are directly linked to the resolution of Standard Model deficiencies.

In light of the above, and considering that HNLs are one of the main motivations for several upcoming intensity-frontier experiments, it appears essential that this potentially complex phenomenology is taken into account by experimental searches, in particular when deriving sensitivity or exclusion limits. To this end, the main objective of this thesis is to investigate several aspects of the phenomenology of HNLs, in relation to current or proposed searches at CERN.

The first article, presented in chapter 2, investigates to which extent the proposed SHiP experiment could probe lepton number violation and HNL oscillations, despite being a beam-dump experiment which a priori should not be able to distinguish lepton number conserving (LNC) and violating (LNV) processes. However, by leveraging the angular distributions resulting from the different spin correlations of LNC and LNV processes, it turns out that SHiP should have a non-trivial sensitivity to lepton number violation, and might even be able to resolve HNL oscillations.

The second article, presented in chapter 3, focuses on a small region of the νMSM parameter space which was recently found to be still allowed [85] by experimental constraints. As an attempt to close this region, a new search channel — consisting of the decay in flight of a K^+ meson into a π^0, a positron e^+ and an HNL N — is investigated at the NA62 experiment. Due to a previouslyunidentified source of background, the sensitivity turns out to be only marginal, unless significant modifications are made to the experiment.

The last article, presented in chapter 4, undertakes to consistently combine the results from an ATLAS search for prompt HNLs [86, 87] with neutrino oscillation data [54, 55] within the νMSM with two nearly-degenerate HNLs. This requires reinterpreting the ATLAS results, and to this end we have implemented a simplified version of their analysis. The reinterpreted limits show a strong dependence on the chosen benchmark point, and often end up above the limit reported for single-flavor mixing. For Dirac-like HNL pairs, non-trivial indirect limits are obtained, where there were previously none. This article emphasizes how limits derived under the single-flavor mixing assumption can greatly differ from limits on the individual mixing angles obtained for non-trivial mixing patterns, thus underscoring the need for reinterpretable results.

In conclusion, the articles collected in this thesis highlight the rich phenomenology of heavy neutral leptons and demonstrate the importance of considering realistic models when performing searches and deriving limits. They show that such models do not necessarily lead to a poorer sensitivity compared to simplified benchmarks, but instead may open new opportunities to probe the underlying theory.

Chapter 2

Dirac vs. Majorana HNLs (and their oscillations) at SHiP

Authors: Jean-Loup Tastet and Inar Timiryasov

Published in *JHEP* **04** (2020), 005

ArXiv: 1912.05520

DOI: 10.1007/JHEP04(2020)005

Abstract

SHiP is a proposed high-intensity beam dump experiment set to operate at the CERN SPS. It is expected to have an unprecedented sensitivity to a variety of models containing feebly interacting particles, such as Heavy Neutral Leptons (HNLs). Two HNLs or more could successfully explain the observed neutrino masses through the seesaw mechanism. If, in addition, they are quasi-degenerate, they could be responsible for the baryon asymmetry of the Universe. Depending on their mass splitting, HNLs can have very different phenomenologies: they can behave as Majorana fermions—with lepton number violating (LNV) signatures, such as same-sign dilepton decays—or as Dirac fermions with only lepton number conserving (LNC) signatures. In this work, we quantitatively demonstrate that LNV processes can be distinguished from LNC ones at SHiP, using only the angular distribution of the HNL decay products. Accounting for spin correlations in the simulation and using boosted decision trees for discrimination, we show that SHiP will be able to distinguish Majorana-like and Dirac-like HNLs in a significant fraction of the currently unconstrained parameter space. If the mass splitting is of order 10^{-6} eV, SHiP could even be capable of resolving HNL oscillations, thus providing a direct measurement of the mass splitting. This analysis highlights the potential of SHiP to not only search for feebly interacting particles, but also perform model selection.

Contents

2.1 Introduction

The experimentally observed non-vanishing neutrino mass differences are among a few firmly established deviations from the Standard Model (SM) predictions. An economic way of generating the light neutrino masses is to introduce heavy singlet fermions with Majorana mass terms into the model [27,30,31,33,34,88]. The masses of the active neutrinos in this extension of the SM are determined by the type-I seesaw formula and at least two singlet fermions are needed to accommodate the two observed mass differences of light neutrinos. A consequence of this mechanism is the presence of heavy Majorana fermions which mix with active neutrinos. The mass scale of these Majorana fermions—Heavy Neutral Leptons (HNLs)—is not fixed. It can be below the electroweak scale,[1]

[1] An argument in favour of the low-scale seesaw comes from the measured values of the Higgs and top masses. HNLs with masses below the electroweak scale are not destabilising the Higgs mass [89,90].

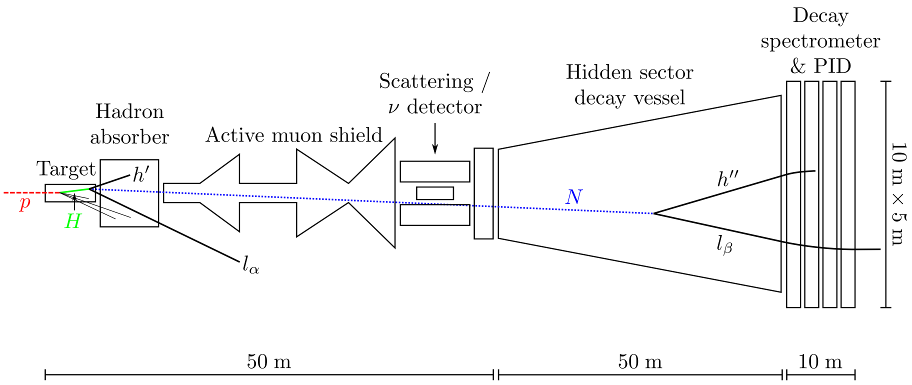

Figure 2.1: Sketch of the SHiP experiment, with the decay chain $H \to h' l_\alpha (N \to l_\beta h'')$.

like in the νMSM [35, 36], where two HNLs are responsible for the light neutrino masses and generating the Baryon Asymmetry of the Universe (BAU) via CP-violating oscillations during their production.

From the FIP (feebly interacting particles) search point of view, HNLs with masses below that of a B meson are the most accessible in the foreseeable future [91]. There is a vast program to search for HNLs at *intensity frontier* experiments, either LHC-based, such as MATHUSLA [41, 92, 93], FASER [94–96], CODEX-b [97, 98], AL3X [43, 99] and ANUBIS [100], or at beam-dump facilities, such as DUNE [101–103] (using the near detector), NA62^{++} [42, 104] (in dump mode) and SHiP [39, 105, 106]. Comparative studies of the *exclusion* limits expected from these experiments have been performed in refs. [107–110]. If a candidate HNL signal were to be observed, the latter three experiments would be sensitive to both its mass and mixing angles.

SHiP is a proposed beam-dump experiment (represented in figure 2.1) set to operate at the CERN SPS. It will use an intense, 400 GeV proton beam from the SPS, dumped on a thick target in order to produce a large number of heavy hadrons, which subsequently decay into Standard Model (SM) or feebly-interacting particles. SHiP is designed to provide a background-free environment to look for the decays of these heavy FIPs. To this end, a hadron absorber located right after the target absorbs most SM particles. It is followed by an active muon shield which deflects the muons away from the experimental cavern. The main detector consists of a decay volume—evacuated in order to reduce the neutrino background, and surrounded by vetos—with a tracker and a calorimeter located at its far end, enabling it to reconstruct the decay event.

In order to generate the light neutrino masses via the seesaw mechanism, HNLs must be Majorana fermions, which violate the total lepton number. However, if the mass splitting is small enough, they can pair to form a coherent superposition of two quasi-degenerate Majorana fermions, which behaves almost like a Dirac fermion. Such a combination is dubbed "*quasi-Dirac pair*". In this case, the mixing angles can exceed the naive seesaw limit $U^2 \approx m_\nu / M_N$ [48, 111, 112], where m_ν and M_N are respectively the mass scales of light neutrinos and HNLs. This is possible because a quasi-Dirac fermion approximately conserves the total lepton number, hence protecting the light neutrino masses. For instance, the νMSM [35, 36] contains such a quasi-Dirac pair if one requires the mass degeneracy which is needed for baryogenesis [35,45] and especially for late-time leptogenesis [82]. Quasi-Dirac pairs also naturally appear in some models of neutrino mass generation, such as the inverse seesaw [113,114] and the linear seesaw [115,116]. This near degeneracy of the HNL masses leads to coherent HNL oscillations. In the νMSM, these oscillations in the early Universe are responsible for baryogenesis.

For sufficiently light ($\lesssim 10\,\mathrm{GeV}$) HNLs like the ones accessible at SHiP, LNV may be experimentally observable even when they form a quasi-Dirac pair [65, 117]. We can distinguish three cases,[2] depending on the scale of the oscillation phase $\delta M\tau$, where δM is the mass splitting of the quasi-Dirac pair and τ the typical proper time probed:

1. **Dirac-like HNL:** One Dirac HNL or a quasi-Dirac pair with an oscillation period exceeding the HNL lifetime or detector size ($\delta M\tau \ll 2\pi$).[3] Only LNC processes can be observed.
2. **Majorana-like HNL:** One Majorana HNL or a quasi-Dirac pair with a lifetime and detector size exceeding the oscillation period ($\delta M\tau \gg 2\pi$). Both LNC and LNV processes can be observed, with equal integrated rates (see section 2.2.2).
3. **Manifestly quasi-Dirac HNLs:** An interesting case occurs when the oscillation period is comparable to the HNL lifetime or to the size of the detector[4] ($\delta M\tau \sim 2\pi$): the experiment may then be sensitive to the coherent oscillations of HNLs.

If HNLs were to be observed at SHiP, the detection or non-observation of lepton number violation and HNL oscillations would allow constraining models and their parameters. The most relevant LNV process at SHiP is the well-studied *same-sign dilepton decay*: $H \to [h'] l_\alpha^+ (N \to h'' l_\beta^+)$, where H, h' and h'' are hadrons (with h' possibly missing), and l_α^+, l_β^+, $\alpha, \beta = e, \mu, \tau$ are charged leptons of potentially different generations. Due to suppressed background, this type of signature is a smoking gun for HNLs in accelerator searches. However, at beam-dump experiments, the heavy hadron decay which produces the HNL takes place inside the target, and therefore the charge of the primary lepton l_α cannot be observed. Naively, it seems that the information about the HNL production is lost, since the charge of the secondary lepton l_β, by itself, is not enough to tell apart LNC and LNV processes. As we shall see in this paper, it turns out that the HNL decay products nevertheless carry important information. Namely, their distribution is different for LNC and LNV processes. Not only does this allow distinguishing Majorana-like from Dirac-like HNLs given sufficiently many events, but the knowledge of these distributions can also be used to *resolve* HNL oscillations and directly measure the mass splitting.

Estimating these two distributions is complicated by the presence of a variety of two- and three-body production channels. In addition, the parent

[2] To be generic, we have included the more exotic cases of a single Dirac or Majorana HNL. The limits presented below are for a quasi-Dirac pair, which only differs from those in the number of events produced.

[3] As pointed out in ref. [117], for most experiments, this possibility might be technically unnatural due to the very small mass splitting needed to satisfy the inequality.

[4] Interestingly, the mass difference needed to generate DM in the νMSM, as found in ref. [82], is exactly in this borderline range.

hadrons are produced with a finite spectrum. As we shall see in section 2.3.3, this smears the distributions, making them look more similar. Therefore, in order to assess whether SHiP will be able to discriminate between Majorana- and Dirac-like HNLs, an accurate treatment of all production channels, including spin correlations, is required. This is accomplished using a Monte-Carlo simulation.

The angular distribution of HNL decay products has been studied in a collider setting for decays which are not fully reconstructible [68,118,119] (such as trilepton decays), as well as for beam-dump experiments [120,121]. Our analysis improves on the latter by not relying on HNLs being produced as helicity eigenstates, by handling a larger class of production channels, by considering the full phase-space distribution of the HNL decay products (instead of just their energy) and by producing a concrete sensitivity estimate using a realistic geometry and heavy meson spectrum for SHiP.

This paper is organized as follows. In section 2.2, we review the Standard Model extended with HNLs, and discuss lepton number violation and coherent HNL oscillations. In section 2.3, we analyze the different signatures of LNC and LNV processes at the SHiP experiment. In section 2.4, we propose a strategy to detect LNV and reconstruct HNL oscillations. Finally, in section 2.5, we present the sensitivity of SHiP to LNV achieved through this method, as well as a possible signature of HNL oscillations. Technical details about the simulation and the statistical analysis are respectively provided in sections 2.A and 2.B.

2.2 Model

2.2.1 Heavy Neutral Leptons

We consider the Standard Model extended with $\mathcal{N}$ HNLs N_I, which are spin-$\frac{1}{2}$ SM singlets with Majorana masses M_I, and new Yukawa couplings $Y^\nu_{\alpha I}$, with $\alpha = e, \mu, \tau$ the lepton flavor index. Using the conventions from [2]:

$$\mathcal{L} = \mathcal{L}_{\mathrm{SM}} + \frac{i}{2} N_I^\dagger (\bar{\sigma} \cdot \partial) N_I - (Y^\nu_{\alpha I})^* (\phi \cdot L_\alpha) N_I - \frac{M_I}{2} N_I N_I + \text{h.c.} \tag{2.1}$$

After electroweak symmetry breaking, the Yukawa interaction generates a Dirac mass term $(m_D)_{\alpha I} = \frac{v}{\sqrt{2}} (Y^\nu_{\alpha I})^*$, resulting in a non-diagonal, symmetric Dirac-Majorana mass term for neutrinos [46]:

$$\mathcal{L}_{\mathrm{DM}} = -\frac{1}{2} \begin{pmatrix} \nu^T & N^T \end{pmatrix} \begin{pmatrix} 0 & m_D^T \\ m_D & M_M \end{pmatrix} \begin{pmatrix} \nu \\ N \end{pmatrix} + \text{h.c.} \tag{2.2}$$

where $M_M = \mathrm{diag}\,(M_I \ldots)$. Using a unitary transformation of the fields (Takagi factorization [47]), the mass matrix can be brought to a diagonal form:

$$\nu_\alpha = U_{\alpha i} n_i \quad \text{and} \quad N_I = U_{Ii} n_i \tag{2.3}$$

$$\mathcal{L}_{\mathrm{DM}} = -\frac{m_i}{2} (n_i n_i + n_i^\dagger n_i^\dagger) \tag{2.4}$$

In the limit $|M_M| \gg |m_D|$, we can use an approximate block factorization, leading to the mass eigenstates $n_i \cong \nu_i, N_I$ mixing with the flavor fields as:

$$\nu_\alpha \cong U^{\mathrm{PMNS}}_{\alpha i} \nu_i + \Theta_{\alpha I} N_I \tag{2.5}$$

$$\Theta_{\alpha I} \cong M_I^{-1} (m_D)_{\alpha I} \tag{2.6}$$

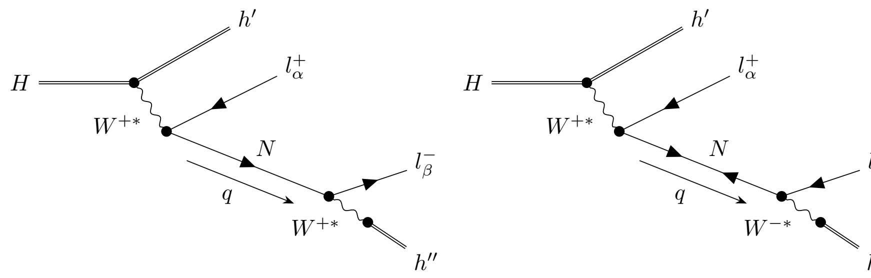

(a) LNC (b) LNV

Figure 2.2: Lepton number conserving and violating decay chains for $H \to h' l_\alpha (N \to l_\beta h'')$.

and the following mass sub-matrices:

$$m_{\alpha\beta} \cong -\sum_I \frac{(m_D)_{\alpha I}(m_D)_{\beta I}}{M_I} \cong -\sum_I M_I \Theta_{\alpha I}\Theta_{\beta I} \tag{2.7}$$

$$m_{IJ} \cong M_I \delta_{IJ} \tag{2.8}$$

The choice of the mass scale M_M and Yukawa couplings $Y^\nu_{\alpha I}$ is not uniquely dictated by low-energy neutrino observables, and should be fixed otherwise.

The Standard Model features an accidental symmetry—lepton number—which, at tree level, is conserved for massless or Dirac neutrinos, but is violated by the Majorana mass term of HNLs. Charged leptons and neutrinos have lepton number $+1$, while charged anti-leptons and anti-neutrinos have lepton number -1. If lepton number is conserved (LNC), then the only allowed Feynman diagrams are those with a conserved flow of lepton number (represented by the arrow on the fermion lines of leptons), like the opposite-sign dilepton decay of a heavy hadron shown in figure 2.2a. On the other hand, in the presence of lepton number violating (LNV) operators, processes like the same-sign dilepton decay shown in figure 2.2b become possible. Lepton number violation can also manifest itself in neutral-current processes or in neutrinoless double-β decay. Whether such LNV transitions actually happen depends on the specific model.

In the past decade, a class of low-scale seesaw models have risen in popularity, such as the νMSM [35], not least because of their falsifiability at existing or proposed experiments. In these models, M_M is postulated to be below the electroweak scale. The seesaw formula (2.7) requires at least 2 HNLs to explain the two observed mass differences. If their parameters are arbitrary, then the smallness of the light neutrino masses is achieved through small Yukawa couplings of order $Y^\nu \sim \frac{1}{v}\sqrt{|m_\nu||M_M|}$, leading to squared mixing angles $|\Theta|^2 \sim |m_\nu|/|M_M|$. For a typical HNL with $M_M \sim 1\,\text{GeV}$, this gives $|\Theta|^2 \sim 10^{-11}$, a number that is too small to be probed at any current or proposed experiment.

However, multiple HNLs can have mixing angles well above the seesaw limit, yet at the same time produce the correct neutrino masses in a technically natural way, if a certain symmetry is imposed on their Yukawa couplings. If we consider for simplicity $\mathcal{N} = 2$ nearly degenerate HNLs $N_{1,2}$, their mixing angles should be related by $\Theta_{\alpha 2} \approx \pm i\Theta_{\alpha 1}$ [48, 111]. Such HNLs form a quasi-Dirac fermion, which approximately conserves the total lepton number. This implies that the usual searches for naive LNV effects (e.g. same-sign dilepton decays), may return null results even if HNLs *are there*.

Below we discuss an important consequence of the approximate nature of this lepton number conservation: HNL oscillations, and how quasi-Dirac HNLs can phenomenologically behave either as Majorana or Dirac HNLs depending on their mass splitting δM and the length scale probed at the experiment.

2.2.2 Coherent oscillations of Heavy Neutral Leptons

The SHiP experiment is only sensitive to GeV-scale HNLs, with mixing angles significantly above the seesaw limit [39]. Therefore it can only probe the quasi-Dirac regime described above. Apart from a small mass splitting $\delta M \ll M$, the two HNLs are otherwise identical. Since these two HNLs cannot be distinguished in any realistic experiment, they both mediate the same processes and each contribute to the total transition amplitude, resulting in interference. Only the initial and final-state particles, which strongly interact with the environment, are measured in the quantum mechanical sense. In order to accurately describe processes involving multiple HNLs, it is therefore necessary to consider them as intermediate particles within a larger process consisting of the HNL production, propagation and decay, and only square the overall transition amplitude between the observed, external particles. This can be formulated rigorously within the framework of the external wave packet model [73, 74] (see also [22, 75–77] and references therein for recent reviews). Let us note in passing that this description automatically takes care of spin correlations between the particles taking part in the HNL production and decay.

In what follows, we consider a typical reconstructible decay chain at SHiP, as depicted in figure 2.2. We will postpone the detailed discussion of this process to section 2.3. A heavy hadron H produced in the target decays at space-time coordinates x_P into an HNL N_I, a charged lepton l_α (the *primary* lepton), and an optional hadron h'. If the HNL is sufficiently long-lived, it can propagate a macroscopic distance before decaying at x_D into a charged lepton l_β (the *secondary* lepton) and a hadron h''.

The slightly different masses of the HNLs mediating the process lead to different dispersion relations $q_I^2 = M_I^2$. As a consequence, the space-time-dependent phase $e^{-iq_I\cdot(x_D-x_P)}$ acquired by the HNL between its production and decay will differ slightly for each mass eigenstate. When squaring the amplitude in order to obtain the differential decay rate, the interference terms between the partial amplitudes coming from different mass eigenstates will therefore feature a space-time-dependent modulation: HNL oscillations. The external wave packet model allows one to unambiguously establish the expression for the oscillation phase and check that the entire process remains coherent in all experimentally relevant situations.

The present paper does not aim to be a detailed study of HNL oscillations, which have already been covered in various settings and limits in the literature [35, 37, 63–66, 68, 69, 80]. Therefore, we will only quote the main result. Let $\mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta}$ be the differential rate for the above-described process $H \to [h']l^\pm_\alpha(N \to l^\pm_\beta h'')$ mediated by a single Majorana HNL N, in the (unphysical) limit of a unit mixing angle between the HNL and the active flavor α at its production vertex, with flavor β at its decay vertex, and without the absorptive part. The coherent differential rate $\mathrm{d}\Gamma^{\pm\pm}_{\alpha\beta}(\tau)$ in the presence of

$\mathcal{N}$ nearly degenerate HNLs mediating the process, as a function of the proper time $\tau = \sqrt{(x_D - x_P)^2}$ between the HNL production and decay vertex, is then:

$$\mathrm{d}\Gamma^{\pm\pm}_{\alpha\beta}(\tau) = \left| \sum_{I=1}^{\mathcal{N}} \Theta^{\pm}_{\alpha I} \Theta^{\pm}_{\beta I} e^{-iM_I \tau - \frac{\Gamma_I}{2}\tau} \right|^2 \mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta} \tag{2.9}$$

where M_I is the (Majorana) mass of the I-th heavy mass eigenstate, Γ_I its total width, and we have used the shorthand notation $\Theta^+ \stackrel{\text{def}}{=} \Theta^*$ and $\Theta^- \stackrel{\text{def}}{=} \Theta$.

In the case of $\mathcal{N} = 2$ HNLs forming a quasi-Dirac pair, i.e. $M_1 = M - \frac{\delta M}{2}$, $M_2 = M + \frac{\delta M}{2}$, $\Theta_{\alpha 2} \cong \pm i \Theta_{\alpha 1}$ and $\Gamma_1 \cong \Gamma_2 \stackrel{\text{def}}{=} \Gamma$, the coherent differential rate reduces to:

$$\mathrm{d}\Gamma^{\pm\pm}_{\alpha\beta}(\tau) \cong 2 \left|\Theta_{\alpha 1}\right|^2 \left|\Theta_{\beta 1}\right|^2 (1 \pm \cos(\delta M \tau))\, e^{-\Gamma\tau} \mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta} \tag{2.10}$$

where the $+$ sign is for lepton number conserving processes ($\mathrm{d}\Gamma^{+-}_{\alpha\beta}$ and $\mathrm{d}\Gamma^{-+}_{\alpha\beta}$), and the $-$ sign for lepton number violating ones ($\mathrm{d}\Gamma^{++}_{\alpha\beta}$ and $\mathrm{d}\Gamma^{--}_{\alpha\beta}$). Notice how in the quasi-Dirac limit, the oscillation pattern does not explicitly depend on the lepton flavors α and β, but only on whether the process is LNC or LNV. If δM vanishes exactly, HNLs form a Dirac fermion and LNV effects are completely absent. Recently, CP-violating HNL oscillations have attracted some interest [70–72, 122]. However, here we can see that CP-violation is suppressed in the quasi-Dirac limit.

Throughout this paper, we will focus on the case where $\Gamma\tau \ll 1$, which is the most relevant for SHiP, and drop the exponentially decaying factor. Analysing formula (2.10), we see that there are three regimes of interest, depending on the mass splitting δM and proper time scale τ probed at the experiment:

- If $\delta M \tau \ll 2\pi$, the HNL pair is observed before the onset of oscillations, and it behaves like a single Dirac HNL, i.e. we cannot observe lepton-number violation.
- If $\delta M \tau \gg 2\pi$, fast oscillations are averaged out, and the HNL pair behaves like a single Majorana HNL, with equal integrated decay rates for LNC and LNV channels.[5]
- If $\delta M \tau \sim 2\pi$, oscillations must be accounted for. If it is possible to experimentally reconstruct, for each selected event, the proper time τ between the production and decay vertex of the HNL, then oscillations can be resolved, i.e. the τ-differential event rates for LNC / LNV will show a periodic modulation according to eq. (2.10).

At SHiP, the proper time scale τ is about 2 m for sufficiently long-lived HNLs. It corresponds to the average time between the production and decay of an *observed* HNL, in its rest frame. Therefore, the critical mass splitting separating the three regimes—near which oscillations are resolvable—is about 10^{-6} eV.

[5]In the rest frame of a single on-shell, Majorana HNL, the only "memory" of the production process is the HNL spin. To perform the phase-space integration for the HNL decay, one can always choose a frame where the HNL is at rest and with a fixed spin projection, hence resulting in the same integrated rates for LNC and LNV processes.

2.3 Probing lepton number violation at SHiP

Many collider searches for Majorana HNLs [123–126] are sensitive to lepton number violation through the charges of the leptons produced at the HNL production and decay vertex. Indeed, due to the chiral nature of the weak interaction, they unambiguously tell the chiral projection through which the HNL interacts at a given vertex. In theory, a same-sign dilepton decay (either prompt or displaced) would thus provide clear evidence for lepton number violation (although, in practice, significant standard model backgrounds exist for prompt decays).

At SHiP, similar numbers of mesons and anti-mesons are expected to be produced.[6] This leads to similar numbers of HNLs being produced along with positively and negatively charged primary leptons. Consequently, the secondary lepton charge contains very little information as to whether the process is LNC or LNV. To lift this degeneracy, it becomes necessary to look at new observables.

Luckily, the HNL lepton number is not the only quantum number conserved by the weak interaction. The HNL also carries spin $\frac{1}{2}$, and the total angular momentum is always conserved. When the HNL is produced, its spin is correlated (opposite if H and h' are pseudoscalar) with that of the primary lepton. Due to chiral suppression, the spin of the primary lepton is itself correlated with its lepton number (see for example the left part of figure 2.3). This suggests that by looking at the angular distribution of the secondary particles—which may be observable—we should be able to obtain information about the primary interaction, and thus whether the process was LNC or LNV (see the right part of figure 2.3). This realization was the starting point of the present work. More generally, we expect LNC and LNV decay chains to have different kinematics due to their different Lorentz structures, potentially allowing us to distinguish them without directly observing the primary decay.

In section 2.3.1, we describe the relevant HNL production and decay channels at SHiP; in section 2.3.2, we quantitatively compare the angular distributions for LNC and LNV processes, and in section 2.3.3 we discuss how this affects the observable momenta in a beam-dump setting.

2.3.1 HNL production and decay at SHiP

At SHiP, most HNLs are produced in heavy meson decays through flavor-changing charged currents, as discussed in ref. [129]. In addition, for the present analysis, we will only consider fully reconstructible HNL decays such as $N \to l_\beta^\mp \pi^\pm$, producing only charged particles which are sufficiently long-lived to be detected by the tracking station located at the end of the decay vessel. Those are also mediated by the charged-current interaction.

Without losing generality, we can therefore consider the generic lepton number conserving and violating processes $H \to [h']l_\alpha(N \to l_\beta h'')$ represented in figures 2.2a and 2.2b, respectively, as well as their CP-conjugates. H denotes a heavy hadron (typically a $D_{[s]}$ or $B_{[c]}$ meson at SHiP), h' and h'' are hadrons

[6] Unless cascade production significantly alters the results from [127]. The charm spectrum will be measured at SHiP prior to data taking [128]. Asymmetries, if present, can only improve the classification accuracy, since the secondary lepton charge would then carry *some* information.

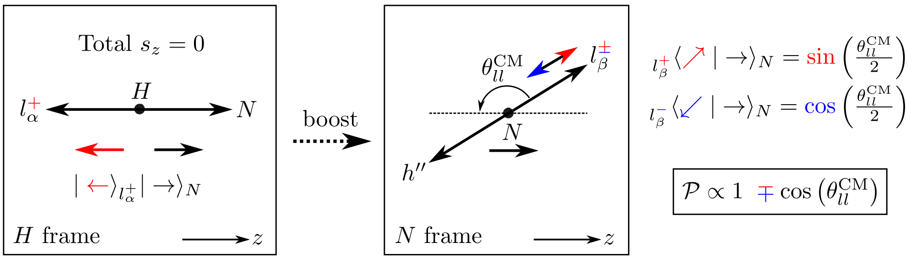

Figure 2.3: This sketch explains the origin of the different angular correlations for LNC and LNV processes. For simplicity, here we consider two-body primary and secondary decays involving only pseudoscalar mesons, and the masses of the charged leptons and of h'' are neglected. For definiteness, the charge of the primary lepton—which is produced inside the target and thus inaccessible—is also fixed to $+$. Since the HNL is a Majorana fermion, the secondary lepton l_β can have either charge. However, due to angular momentum conservation, the lepton l_α^+ and the HNL N are produced with opposite spin projections in the rest frame of the heavy meson H. Because of chiral suppression (which is more effective for light fermions), the charge of the primary lepton is correlated with its spin (e.g. in the massless limit, l_α^+ has helicity $+\frac{1}{2}$) and hence with the HNL spin. For the same reason, the angular distribution of the decay products of the resulting HNL spin eigenstate (which is unaffected by a boost along the quantization axis) will therefore depend on the secondary lepton charge. The very same formula for the probability $\mathcal{P}$ also holds for CP-conjugated channels, with the $+$ sign for LNC and the $-$ sign for LNV. The general case (massive, with two- or three-body primary decay) is discussed in section 2.3.2.

(with h' missing for two-body primary decays), and $l_\alpha^\pm$ and $l_\beta^\pm$ are respectively the primary and secondary leptons.

Since the heavy hadron H is typically short-lived, the primary decay takes place inside the target and cannot be observed. If the HNL is sufficiently long-lived (we will assume this to be the case throughout this paper), it can propagate a macroscopic distance before decaying, and leave a very displaced vertex inside the SHiP decay vessel. For the selected decay channels $N \to l_\beta^\mp \pi^\pm$, this secondary vertex can be fully reconstructed.

In the present study, we will restrict ourselves to HNL masses between the K and D_s thresholds. Masses below the K threshold have already been heavily constrained [91], while above the D_s mass, HNLs are mainly produced in B meson decays, whose spectrum cannot be directly measured at the beam dump, making our analysis more sensitive to modeling errors.

2.3.2 Angular correlations in LNC and LNV decay chains

In order to study the angular correlations between all final-state particles, spin correlations between the primary and secondary decay must be accounted for. Those result from the non-observation of the HNL spin, which leads to interference between the two spin eigenstates N_s, $s = \pm\frac{1}{2}$ (similarly to how the non-observation of its precise mass allows for flavor oscillations). To compute the overall transition amplitude, we can therefore use the same trick as for oscillations, i.e. treat the primary and secondary decays as a single process.

To simplify the calculations, in this section we will focus on the case of a single Majorana HNL, which mediates both LNC and LNV decay chains with equal rates, and we will omit the absorptive part of the amplitude (i.e. we will study $\mathrm{d}\hat{\Gamma}^{\pm\pm}_{\alpha\beta}$ instead of $\mathrm{d}\Gamma^{\pm\pm}_{\alpha\beta}(\tau)$). We do not lose generality in doing so, because the effect of multiple nearly degenerate HNLs and their finite lifetime can be factored out, and subsequently recovered, using eqs. (2.9) and (2.10). To keep the notation light, we will from now on drop the HNL index $I = 1$.

Since we are only concerned with long-lived HNLs, which are produced on their mass shell and have well separated, localized production and decay vertices, the momentum q of the HNL is practically fixed, which allows factorizing the transition amplitude as:

$$\mathcal{A}\left(H \to h' l_\alpha l_\beta h''\right)\Big|_{N \text{ long-lived}} \propto \sum_{s=\pm\frac{1}{2}} \mathcal{A}\left(H \to h' l_\alpha N_s(q)\right) \mathcal{A}\left(N_s(q) \to l_\beta h''\right) \tag{2.11}$$

where we have omitted the complex phase $e^{-iq\cdot(x_D - x_P)}$ resulting from the HNL propagation, which is unimportant in the case of one HNL. The sub-amplitudes for the primary and secondary polarized decays are then straightforward to compute using the usual Feynman rules with two-component spinors [2].

Consider now the LNC and LNV processes $H \to [h'] l_\alpha (N \to l_\beta h'')$ where H, h', h'' are pseudoscalar mesons and h' may be missing. They are respectively represented in figures 2.2a and 2.2b, with the arrows denoting the flow of lepton number. Their CP-conjugates have been omitted, since in the absence of oscillations (as is the case for the incoherent width), CP is conserved. As can be seen in figure 2.5, the primary decays $H \to [h'] l_\alpha N$ with h' a pseudoscalar meson or missing indeed produce the majority of HNLs with masses $\gtrsim 0.7\,\mathrm{GeV}$ and below the D_s mass.[7] Let $J^h_{W\mu}$ be the hadronic charge-lowering current, $j^-_{1\mu} = \langle big\rangle h' | J^h_{W\mu} | H$ and $j^\mp_{2\mu} = \langle big\rangle h'' | J^{h(\dagger)}_{W\mu} | 0$ the hadronic matrix elements, $p_{\alpha,\beta}$ the charged lepton momenta, and q the HNL momentum. If the primary decay is purely leptonic, then $|h'\rangle = |0\rangle$. Since SHiP cannot directly measure the spin or helicity of the particles detected, we sum incoherently over all possible spin configurations of final state particles. The spin-summed, squared amplitudes are then, in the Fermi approximation:

$$\overline{\left|\mathcal{A}_{\mathrm{LNC}}(H \to h' l^+_\alpha l^-_\beta h'')\right|^2} = \frac{\left|\Theta_\alpha\right|^2 \left|\Theta_\beta\right|^2}{v^8} \operatorname{tr}\left(P_R \not{p}_\alpha \not{j}^*_1 \not{q} \not{j}^*_2 \not{p}_\beta \not{j}_2 \not{q} \not{j}_1\right) \tag{2.12}$$

$$\overline{\left|\mathcal{A}_{\mathrm{LNV}}(H \to h' l^+_\alpha l^+_\beta h'')\right|^2} = \frac{\left|\Theta_\alpha\right|^2 \left|\Theta_\beta\right|^2}{v^8} M_N^2 \operatorname{tr}\left(P_R \not{p}_\alpha \not{j}^*_1 \not{j}^*_2 \not{p}_\beta \not{j}_2 \not{j}_1\right) \tag{2.13}$$

where we have omitted the $^\pm$ for brevity if they can be inferred from context, $\Theta_{\alpha,\beta}$ are the mixing angles, and $v = \langle|\phi|\rangle \approx 246\,\mathrm{GeV}$ is the vacuum expectation value of the Higgs field. These results are consistent with the polarized decay rates from [103], but generalize to the case where the primary decay produces a superposition of HNL helicity eigenstates. The above two expressions differ in

[7] Below $M_N \approx 0.7\,\mathrm{GeV}$, a non-negligible fraction of HNLs is produced along with a *vector* meson. In this case, we expect the angular correlations to reverse compared to the pseudoscalar case.

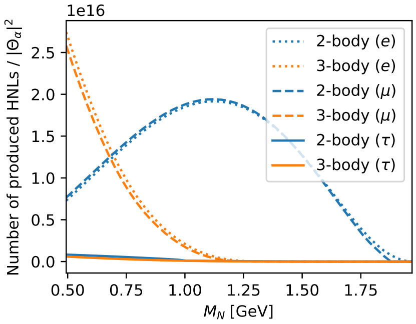

Figure 2.4: Number of HNLs produced at SHiP as a function of the primary decay multiplicity, for a coupling to one flavor.

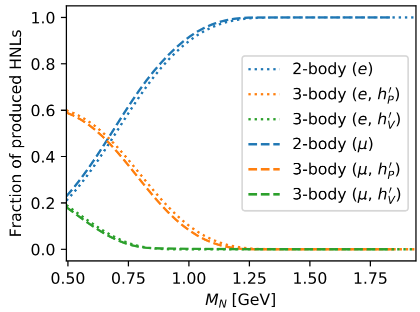

Figure 2.5: Fraction of HNLs produced at SHiP as a function of the primary decay multiplicity and spin of the outgoing meson, for a coupling to one flavor.

the trace, therefore we generically expect them to produce different momentum distributions for LNC and LNV processes. However, in their current form, this difference is not manifest. To understand it, it is interesting to consider the special case where the production process is a two-body decay. As can be seen in figures 2.4 and 2.5, it is actually the main production channel for HNLs with masses $\gtrsim 1\,\mathrm{GeV}$ and below the D_s mass.

When both the production and decay process are two-body decays, the hadronic matrix elements are $j_1^\mu = -iV_{UD} f_H p_H^\mu$ and $j_2^\mu = +iV_{U'D'} f_{h''} p_{h''}^\mu$, where V_{UD} denotes the relevant CKM matrix element and f_h is the meson decay constant. Neglecting the masses of the final state particles, which give $\mathcal{O}\left(m^2_{\alpha,\beta,h''}/M^2_{H,N}\right)$ corrections, the traces from eqs. (2.12) and (2.13), respectively for LNC and LNV processes, simplify to:

$$\mathrm{tr}\left(P_R \not{p}_\alpha \not{j}_1^* \not{q} \not{j}_2^* \not{p}_\beta \not{j}_2 \not{q} \not{j}_1\right) \cong |V_{UD}|^2 \, |V_{U'D'}|^2 \, f_H^2 f_{h''}^2 \cdot M_N^6 \left(M_H^2 - M_N^2 - s_{ll}\right) \tag{2.14}$$

$$M_N^2 \mathrm{tr}\left(P_R \not{p}_\alpha \not{j}_1^* \not{j}_2^* \not{p}_\beta \not{j}_2 \not{j}_1\right) \cong |V_{UD}|^2 \, |V_{U'D'}|^2 \, f_H^2 f_{h''}^2 \cdot M_N^6 s_{ll} \tag{2.15}$$

where $s_{ll} \overset{\text{def}}{=} (p_\alpha + p_\beta)^2$ is the invariant dilepton mass. Note the linear and opposite dependences of the LNC and LNV spin-summed squared amplitudes on s_{ll}. To understand their origin, it is enlightening to reexpress s_{ll} in the rest frame of the HNL, in terms of the angle $\theta_{ll}^{\mathrm{CM}} = \angle(\mathbf{p}_\alpha^{\mathrm{CM}}, \mathbf{p}_\beta^{\mathrm{CM}})$ between the two lepton momenta. Still in the massless limit, we find:

$$s_{ll} = \frac{M_H^2 - M_N^2}{2}\left(1 - \cos\left(\theta_{ll}^{\mathrm{CM}}\right)\right) \tag{2.16}$$

Therefore,

$$\overline{|\mathcal{A}_{\mathrm{LNC}}|^2} \propto 1 + \cos\left(\theta_{ll}^{\mathrm{CM}}\right) \tag{2.17}$$

$$\overline{|\mathcal{A}_{\mathrm{LNV}}|^2} \propto 1 - \cos\left(\theta_{ll}^{\mathrm{CM}}\right) \tag{2.18}$$

We observe that opposite-sign leptons (LNC) tend to be produced in the same direction, and same-sign leptons (LNV) in opposite directions. As explained in figure 2.3, this is a consequence of the chirality of the weak interaction and the conservation of the total angular momentum. In the absence of any other dynamics, spin projections lead to the characteristic angular dependence in $\cos\left(\theta_{ll}^{\mathrm{CM}}/2\right)$ and $\sin\left(\theta_{ll}^{\mathrm{CM}}/2\right)$ of the transition amplitude, respectively for LNC and LNV. Equations (2.17) and (2.18) then directly follow from squaring the amplitude.

In the massive case, the finite masses of the decay products can result in helicity flips, and in the three-body case, the QCD matrix elements lead to non-trivial correlations between the momenta of the primary decay products. These effects complicate the correlations between the various momenta. Nevertheless, they can be accounted for when *sampling* events. To this end, we have implemented the full matrix elements from eqs. (2.12) and (2.13) in our Monte-Carlo simulation, as discussed in section 2.A.4.

2.3.3 Angular distribution in the laboratory frame

At SHiP, the invariant mass s_{ll} (or angle $\theta_{ll}^{\mathrm{CM}}$) cannot be reconstructed. This is because neither the heavy hadron momentum nor the momenta of its decay products (other than the HNL) can be determined. Indeed, the heavy hadrons producing the HNLs do not have a monochromatic spectrum, and the primary decay cannot be observed since it takes place inside the target. One can then reasonably wonder if some difference between the LNC and LNV distributions subsists when looking only at the (observable) secondary decay products, in the laboratory frame, or if it is washed out.

To start answering this question, it is instructive to go back to the simplified case discussed in section 2.3.2, where the HNL is produced and decays through two-body processes involving pseudoscalar mesons. In the HNL rest frame, we obtained the following correlation: for LNV processes, the direction of the secondary lepton momentum is positively correlated with the boost direction (denoted by z on figures 2.3 and 2.6) from the heavy meson rest frame to the HNL rest frame; while for LNC processes it is anti-correlated. This is depicted in the left panel of figure 2.6. Furthermore, in two-body decays, the magnitudes of all momenta in the rest frame of the parent particle are fixed by four-momentum conservation, and depend only on the particle masses. Consequently, in the heavy meson rest frame, the magnitude of the secondary lepton momentum will on average be larger for LNV processes compared to LNC ones. This argument is still valid for three-body decays involving pseudoscalar mesons. A non-trivial asymmetry thus subsists in the heavy meson rest frame (see the middle panel of figure 2.6).

As a final step, the momenta must be boosted back to the laboratory frame. Since the heavy hadron momentum is not fixed, this has the potential to wash out the correlations. At SHiP, heavy mesons have a large momentum spread along the beam axis ($\mathcal{O}(10\,\mathrm{GeV})$, much larger than the yield of the meson decay), and a significantly smaller one ($\mathcal{O}(1\,\mathrm{GeV})$) in the transverse direction (see section 2.A.3). The asymmetry between the LNC and LNV distributions is therefore more likely to be visible in the transverse plane than along the beam axis. For it to be significant, the HNL kinetic energy in the heavy hadron rest frame should be similar to or exceed the transverse momentum spread of the

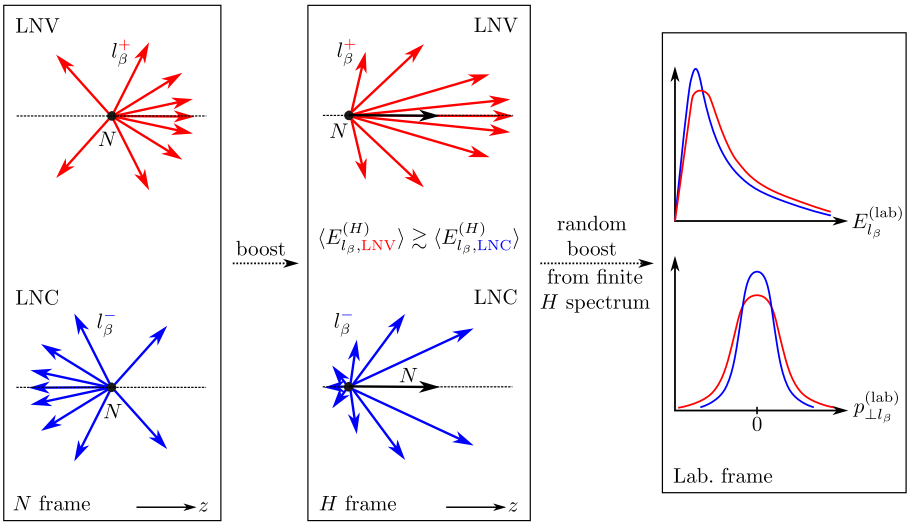

Figure 2.6: This sketch shows how the different distributions of l_β in the HNL rest frame for LNC vs. LNV processes affect the corresponding distributions in the rest frame of the heavy hadron H and in the laboratory frame. The various momenta shown for l_β represent multiple realizations of the decay. In the H frame, LNV processes typically result in larger momenta for l_β than LNC ones. In the laboratory frame, this effect partly survives the averaging over the heavy hadron spectrum and manifests itself as a broadening of the distribution of the secondary lepton momentum p_β.

hadron spectrum. As a result, we expect the p_T spectrum of the secondary lepton l_β to be broader for LNV processes than for LNC ones (see the right panel of figure 2.6), provided that both of them are broader than the irreducible p_T spread of the heavy meson spectrum.

Alternatively, one could try to approximate the angle $\theta_{ll}^{\mathrm{CM}}$ in the HNL rest frame. If the heavy hadron momentum is fixed, this can be done exactly, and results in the maximal classification accuracy allowed by spin projections (e.g. $a = 3/4$ in the two-body, massless case). It is then equivalent to measuring the (observable) momentum p^{CM} of the secondary lepton l_β in the HNL rest frame. However, when the heavy hadron has a finite spectrum, the boost direction from its rest frame to the HNL rest frame is not fixed any more. This partially decorrelates $\theta_{ll}^{\mathrm{CM}}$ and p^{CM}, hence reducing the discriminating power of the latter.

As we shall see in section 2.4.2, the features discussed above can indeed be used to discriminate between LNC and LNV processes (see for example figure 2.7). More generally, any difference—in the laboratory frame—between the distributions of the visible decay products of LNC and LNV processes opens up the possibility of measuring their relative rates, given sufficiently many events. Although discriminating between these two classes of events would be very challenging analytically, this problem is well suited to multivariate analysis.

Further complications arise, however, due to HNLs being produced from a mix of various two- and three-body decays, and because of the geometrical acceptance of the experiment, which alters the distribution of visible particles. Generating a training set which faithfully reproduces the angular correlations discussed above while including these effects is therefore best done using a Monte-Carlo simulation. In the next section, we discuss the simulation used to generate the training set (section 2.4.1), then how we use it to train a binary classifier (section 2.4.2), and finally how we use the classifier output in order to perform model selection (section 2.4.3) and reconstruct HNL oscillations (section 2.4.4). In section 2.4.5, we discuss the applicability of the method presented here to other proposed experiments.

2.4 Simulation and analysis

2.4.1 Simulation

In order to accurately estimate the distribution of the momenta of the HNL decay products, we have devised a simple Monte-Carlo simulation, which generates the primary and secondary decays at once, using the matrix elements presented in section 2.3.2. The first step is to generate D mesons with a realistic spectrum. Generating these spectra from simulation would be a difficult undertaking, so instead we chose to use experimental data collected by the LEBC-EHS collaboration [127], at the CERN SPS running at 400 GeV with a hydrogen target. We then randomly select a production and decay channel according to the relative abundances of charmed mesons from [106] and the branching fractions from [129]. Finally, we generate the momenta of both the primary and secondary decay products at once. This is done by first sampling all the momenta according to phase-space, independently for each decay, and finally performing rejection sampling on these momenta using the matrix element for the combined process. As a last step, we simulate the geometrical acceptance by requiring the HNL to decay within the hidden sector decay vessel, into two long-lived, charged particles which both intersect the tracking station. In order to account for the (small) probability of the HNL decaying inside the fiducial volume, each event is weighted by $P_{\text{decay}}(\tau) = \Gamma e^{-\Gamma\tau}$, where τ is the proper time between the HNL production and decay. Throughout this paper, we assume the particle identification to be perfectly efficient, which should be a reasonably good approximation at SHiP [130]. The simulation is described in details in section 2.A.

2.4.2 LNC / LNV classification

For a given choice of relative squared mixing angles $|\Theta_\alpha|^2$ (which are supposed to be known by the time LNV is studied at SHiP), we generate a dataset for a range of HNL masses between the K and D_s thresholds. For each HNL mass, we sample $3 \cdot 10^6$ events with uniform weights, and keep only those passing the acceptance cuts. The HNL is simulated as a single Majorana particle, which ensures that the dataset contains equal numbers of LNC and LNV events, and is also balanced with respect to the primary and secondary lepton charges.

Each event is labelled with a boolean flag set to false for LNC and true for LNV, using the MC truth. The only observable quantities come from the HNL decay in the vacuum vessel. They are: the momenta and charges of the lepton $l_\beta^\pm$ and pion $\pi^\mp$, and the decay vertex x_D. Of these quantities, we

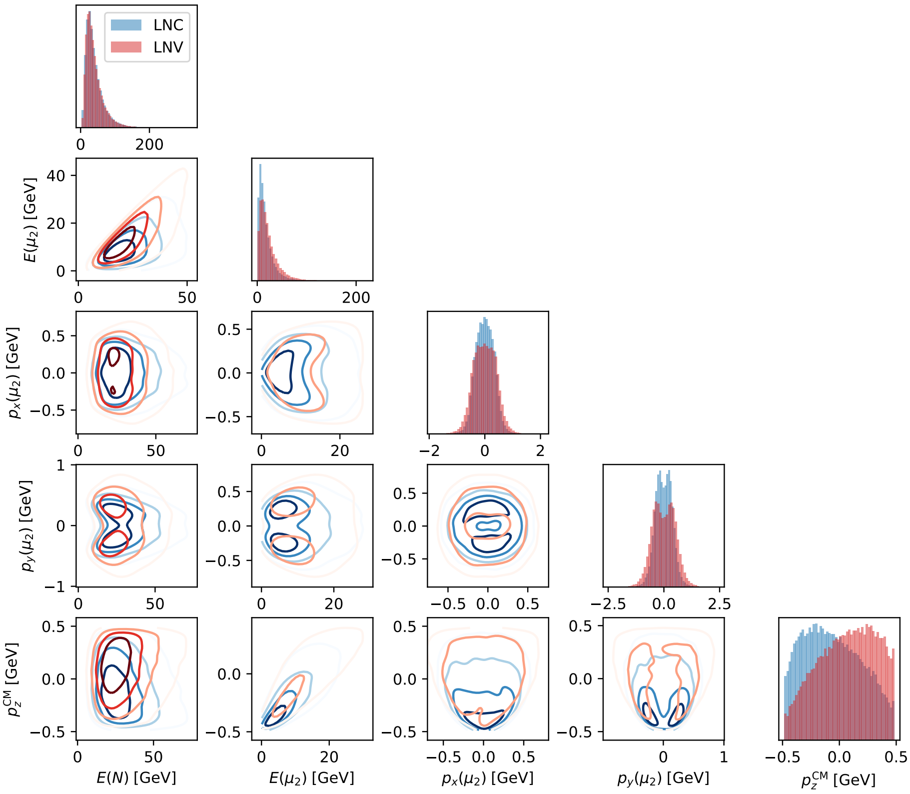

Figure 2.7: Corner plot showing the correlations between five selected features, for a 1 GeV HNL coupling to the muon. See table 2.1 for a description of the features. Each subplot shows, on the same scale, the marginal distributions of LNC and LNV events as a function of either one (on-diagonal plots) or two (off-diagonal plots) features. 1d distributions are represented as histograms, and 2d distributions as contour plots of the probability density.

record a total of 19 primary or derived features. Their definitions can be found in table 2.1, and some typical distributions are presented, as an example, in figure 2.7, for both LNC and LNV processes. Finally, from each dataset, we set aside 30% of events for testing and 20% for validation, leaving us with 50% of events for training the classifier.

For each dataset, we train a binary classifier to discriminate between LNC and LNV decay chains. For this study, we use the LightGBM [131] decision tree boosting algorithm, through the Python interface to the reference implementation [132]. In order to perform simple classification, we choose the `binary` objective. The training is discussed in more details in section 2.B.2. The accuracy of the trained classifier (as evaluated on the test set) is presented in figure 2.8 as a function of the HNL mass for three scenarios, corresponding to an HNL coupling to electrons, muons, or equally to both.

Feature(s)	Description
`Q12`	Charge of the secondary lepton l_β
`E1, p1x, p1y, p1z`	Reconstructed HNL momentum $p_N = p_{l_\beta} + p_\pi$ (lab frame)
`E2, p2x, p2y, p2z`	Secondary lepton momentum p_{l_β} (lab frame)
`E3, p3x, p3y, p3z`	Secondary pion momentum p_π (lab frame)
`pCMx, pCMy, pCMz`	Secondary lepton momentum p_{CM} (HNL frame)
`xD, yD, zD`	Decay vertex (lab frame)

Table 2.1: The 19 features recorded for each event.

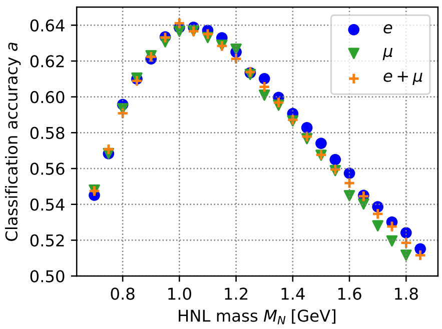

Figure 2.8: Classification accuracy as a function of the mass, for an HNL coupling to e, μ, or equally to both.

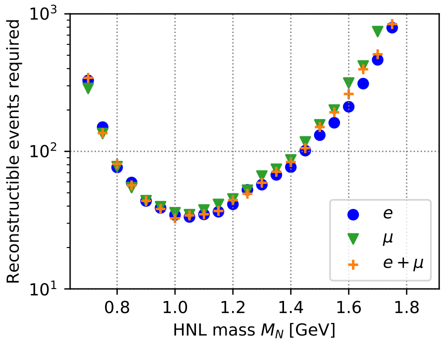

Figure 2.9: Number of fully reconstructible events required to detect LNV at 90% CL, for an HNL coupling to e, μ, or equally to both.

2.4.3 Model selection

Assuming the true event distribution to match (or be sufficiently close to) the simulated one, we can then use our trained classifier to classify each event as either LNC or LNV. As stated in section 2.1, our main goal is to distinguish the following two hypotheses:

- $\mathcal{H}_1$: HNLs are Dirac or quasi-Dirac with $\delta M \tau \ll 1$ (LNC decays only).
- $\mathcal{H}_2$: HNLs are Majorana or quasi-Dirac with $\delta M \tau \gg 1$ (as many LNC / LNV decays).

Since the classifier is not perfectly accurate, its decision cannot be used to directly confirm the presence of LNV processes, or constrain their existence. If we knew the full distribution in feature space $\rho(z)$ for each hypothesis, we could obtain an optimal test statistics by constructing the corresponding likelihood ratio [133]. However, accurately estimating $\rho(z)$ is a non-trivial task and would be error-prone, so we elected to use a less powerful but more reliable, simplified model. Knowing the classification accuracy a for a given binary classifier, we compute the likelihood of classifying k events out of N as LNV, and $N-k$ events as LNC (independently of their specific feature vectors z) assuming that the true fraction of LNV events is f. We then compute the best-fit value for f and use Wilk's theorem [134] in order to determine whether it significantly deviates from either $f = 0$ (corresponding to $\mathcal{H}_1$) or $f = \frac{1}{2}$ (corresponding to $\mathcal{H}_2$).

In order to estimate the "model-selection" sensitivity of SHiP, we then compute, under each hypothesis and as a function of the HNL mass M_N and squared mixing angles $|\Theta_\alpha|^2$, the median confidence level at which we can exclude the other hypothesis assuming 5 years of nominal operation (i.e. $2 \cdot 10^{20}$ protons on target). For each true hypothesis, we finally draw the sensitivity limit by plotting, for each M_N, the smallest $|\Theta_\alpha|^2$ for which this median confidence level is at least 0.9. In other words, for mixing angles above this limit, SHiP has a probability of at least 1/2 of disfavouring one hypothesis at CL = 0.9 if the other is realized. The number of fully reconstructible events corresponding to this limit is plotted in figure 2.9 (when the null hypothesis is taken to be $\mathcal{H}_1$). The construction of these confidence limits is described in details in section 2.B.3, and the resulting sensitivity plots are presented in section 2.5.1.

2.4.4 Resolving HNL oscillations

So far we have only considered the two extreme cases ($\mathcal{H}_1$ and $\mathcal{H}_2$), where the HNL(s) behave either as a single Dirac or Majorana particle. However, as discussed in section 2.2.2, if two nearly degenerate HNLs form a quasi-Dirac pair, both LNC and LNV decay chains will be present, with a non-trivial ratio $\neq 0, 1$, and the corresponding decay rates will feature oscillations as a function of the proper time τ between the HNL production and decay events, with the characteristic $1 \pm \cos(\delta M \tau)$ dependence described by eq. (2.10), where $(+)$ corresponds to LNC and $(-)$ to LNV.

For $\delta M \sim 10^{-6}$ eV, $\delta M \tau$ will be of order 2π at SHiP, leading to potentially resolvable oscillations, provided we can accurately reconstruct the proper time τ between the HNL production and decay. Expressing it as $\tau = {}^{L}/_{\beta\gamma}$, we see that this can be accomplished if we have sufficiently accurate vertexing and energy reconstruction. At SHiP, the precision on L will be limited by the impossibility of reconstructing the primary vertex within the target. The energy resolution, despite being sufficient for particle identification, is not enough for reconstructing τ (see sections 4.7 and 4.10 in ref. [105]). However, the momentum resolution, combined with the dispersion relation (assuming the HNL mass to be known already with sufficient accuracy) should allow reconstructing γ much more precisely. The high vertexing and momentum resolution permitted by the SHiP tracker, together with our method for (statistically) distinguishing LNC from LNV processes (described in section 2.4.3), should therefore make it possible to resolve the oscillation pattern in part of the parameter space.

In order to search for HNL oscillations, we first classify the observed events using a model trained (for *one* HNL) at the corresponding mass. We thus assume again that we have sufficiently many events that the HNL mass M_N is well known. The events are then binned in proper time τ, which is the relevant variable for oscillations of massive, relativistic particles. Instead of using the predicted class, here we implement the classifier decision as a weight for the binned events, using the predicted probability p_{LNV}. This weight contains more information than the class does, since it acts as a measure of uncertainty by taking values close to 1/2 for ambiguous events, and closer to 0 or 1 for unambiguous ones. However, without applying further corrections, the sum of these probabilities would average to $N \langle p_{\mathrm{LNV}} \rangle$ for the entire sample of N events. If used directly as weights, they would therefore cause the oscillatory pattern to be hidden among Poisson fluctuations. In order to reveal this pattern, we instead weight the events by $p_{\mathrm{LNV}} - \overline{p_{\mathrm{LNV}}}$, where $\overline{p_{\mathrm{LNV}}}$ is the sample average of the estimated p_{LNV}. This weight averages to zero over the entire sample, which limits the impact of Poisson fluctuations.

HNL oscillations are implemented in our simulation by first generating events without taking interference into account then, in a second time, performing rejection sampling based on the proper time τ, following eq. (2.10). The results obtained using this simulated data set are presented in section 2.5.2.

2.4.5 Applicability of the method to other experiments

The present study crucially relies on the identification of the HNL decay products and the measurement of their momenta. However, a number of proposed experiments to search for HNLs, such as MATHUSLA [41, 92, 93], CODEX-b [97, 98] (in its baseline configuration) and ANUBIS [100], cannot measure the momenta of the decay products. Since low-mass HNLs ($M_N < M_{B_c}$) at the LHC are also mostly produced in the decays of heavy mesons, one can wonder to which extent the present analysis would apply to these experiments. Training a classifier using only the *directions* of the tracks of the visible decay products and the same geometry as SHiP reveals that the distributions of LNC / LNV for a given set of HNL parameters can still be distinguished, with an accuracy only slightly lower than the one obtained using the full momenta. There are, however, two caveats. First, training the classifier requires knowing the HNL mass, which cannot be obtained without measuring the momenta of its decay products (or matching the displaced decay to its reconstructed production process in the main detector, if this is feasible). In addition, the large center-of-mass energy at the LHC could result is a very broad heavy meson spectrum, which would smear out the LNC / LNV distributions and make them indistinguishable. It therefore seems unlikely that MATHUSLA, CODEX-b or ANUBIS could benefit from this method.

Other planned or proposed detectors, such as NA62^{++} [42, 104] (in beam-dump mode), the DUNE near detector [101–103], FASER [94–96] and AL3X [43, 99], are in principle capable of reconstructing the HNL mass. The AL3X detector, thanks to its large time projection chamber and its magnetic field, should be able to directly measure both the charges and momenta of the two leptons, making the method described here unnecessary. It is unclear to the authors, however, whether FASER could benefit from it. The answer likely depends on the spectrum of the heavy mesons producing the HNLs which eventually interact with the detector. A Monte-Carlo simulation would provide a definitive answer to this question. The remaining beam-dump experiments: NA62^{++} and DUNE, share a similar geometry with SHiP and face the same challenge (no observation of the primary charged lepton $l_\alpha^\pm$). As such, we generically expect the method presented here to be applicable to these experiments, within the mass range where it is valid, and subject to the heavy meson spectrum being similar to the one at SHiP. This could be ascertained using a Monte-Carlo simulation. Whether these experiments can also resolve HNL oscillations will depend on how accurately they can reconstruct the HNL momentum.

2.5 Results

2.5.1 Sensitivity to Lepton Number Violation

In order to easily compare our results to existing exclusion bounds or to the sensitivities of future experiments, let us consider two simplified models where a single HNL exclusively mixes with the electron or muon neutrino.[8] As can be seen in figure 2.8, more generic mixing patterns with the e and μ flavors do not significantly degrade the classification accuracy; therefore they should leave the limits presented below mostly unchanged. However, if a significant fraction of HNLs is produced through mixing with the τ neutrino, then the present analysis would need to be modified to handle secondary production of HNLs in τ decays, including spin correlation effects.

As discussed in section 2.4.3, we define the sensitivity to lepton number violation as the smallest mixing angles for which SHiP has a 1/2 probability of either rejecting or detecting LNV, if it is respectively absent or present with the same rate as LNC. The results are presented in figure 2.10, along with various existing exclusion bounds and detection sensitivity[9] limits for planned or proposed experiments, extracted from the report of the Physics Beyond Colliders working group [91]. We only show the sensitivities of experiments which can not only set exclusion bounds, but also reconstruct the HNL mass, should it be observed. Note that in order to be consistent with the SHiP detection sensitivity, which was computed for one Majorana HNL, we present our results for one HNL as well. In the realistic case of $\mathcal{N} \geq 2$ HNLs, both curves must be scaled down by a factor of $\mathcal{N}^{1/2}$. Above the black dashed line, SHiP should be able to distinguish Dirac-like ($\mathcal{H}_1$) and Majorana-like ($\mathcal{H}_2$) HNLs. We have discarded the HNL masses for which the early stopping criterion returned the first iteration as the best, since it suggests that the classifier has failed to learn anything about the data. Below 0.7 GeV, additional production channels $H \to h'_V l_\alpha N$ (where h'_V denotes a vector meson) become significant, and have not been implemented with spin correlations in our Monte-Carlo simulation. Therefore we also restrict the HNL mass to $M_N \gtrsim 0.7\,\text{GeV}$. Additionally, since the sensitivity is almost identical for excluding $\mathcal{H}_1$ or $\mathcal{H}_2$, we only plot one limit, which corresponds to excluding $\mathcal{H}_1$ at 90% CL if LNV is actually present.

We can see that the larger number of accepted events (indicated in figure 2.10 by the thin dashed grey lines) at higher masses initially compensates for the worse classification accuracy, but is not sufficient any more as we approach the D threshold. In practice, we expect that systematic uncertainties about the D spectrum and the simulation will decrease the sensitivity at both ends of the mass range, where the classification accuracy is already close to 1/2. Comparing the results to the SHiP detection sensitivity, we see that around 1 GeV, the model-selection sensitivity limit is about one order of magnitude above the detection one, while remaining well below the planned NA62^{++} limit as well as existing bounds.

[8] Within the seesaw mechanism, it is impossible to generate the two observed light neutrino mass differences with a single HNL, or if HNLs mix with one generation only [135]. The two benchmarks presented in figures 2.10a and 2.10b are thus simplifications, used here because they are consistent with the parametrization employed by the PBC working group.

[9] The usual sensitivity, by opposition to the sensitivity to lepton number violation discussed here.

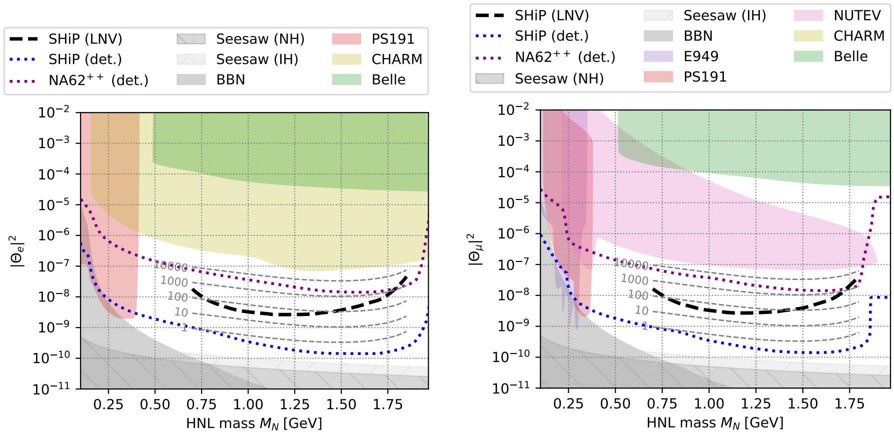

(a) HNL mixing with ν_e. (b) HNL mixing with ν_μ.

Figure 2.10: SHiP sensitivity to lepton number violation. The thick dashed curve is the "model-selection" sensitivity computed in this work. The thin dashed grey lines show the number of fully reconstructible events which would be observed at SHiP for a given mass and mixing angle. Dotted curves are the (lower) detection sensitivities for the proposed or planned experiments which can reconstruct the HNL mass. Coloured, filled areas are regions of parameter space which have been excluded by previous experiments. The grey filled area denoted by BBN indicates the region which is incompatible with Big Bang Nucleosynthesis. Below the seesaw limit[10](hatched region), mixing angles are too small to produce the observed neutrino masses.

This leads us to an interesting conclusion: there exists a non-trivial region of parameter space, unconstrained by current or near-future experiments, where SHiP would not only be able to detect HNLs, but also characterize them as either Dirac-like or Majorana-like particles. As discussed in sections 2.A.3 and 2.B.4, this conclusion is robust with respect to uncertainties on the heavy meson spectrum.

2.5.2 Resolvable quasi-Dirac oscillations

The result of the procedure described in section 2.4.4 is presented in figure 2.11 for a new simulated dataset (independent from the training set), corresponding to a quasi-Dirac pair of mass $M_N = 1\,\text{GeV}$, mass splitting $\delta M = 4 \cdot 10^{-7}\,\text{eV}$, and mixing with muon neutrinos only, with a squared mixing angle $\left|\Theta_{\mu I}\right|^2 = 2 \cdot 10^{-8}$, $I = 1, 2$. The oscillatory pattern is manifest at $\tau < 5\,\text{m}$, where most of the events fall. At larger τ it is hidden in Poisson fluctuations. The uncertainty on τ at SHiP is dominated by the (boosted) length of the target $\sim 0.1\,\text{m}$, which contains the unresolved primary vertex. It could smear out fast oscillations,

[10]The seesaw limit can only be rigorously computed if the mixing angles are consistent with the seesaw equation (2.7). This is not possible for HNLs mixing with only one generation, nor for a single HNL. The limits presented here instead correspond to the "naive" estimate $\sum m_\nu \leq M_N \cdot \sum_\alpha |\Theta_\alpha|^2$, where we have assumed the lightest neutrino to be massless.

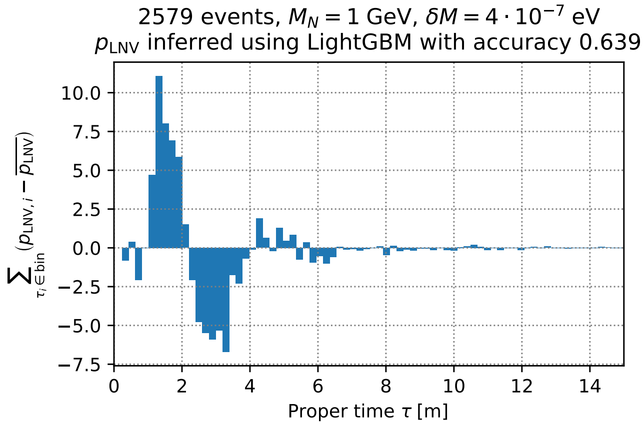

Figure 2.11: Events binned by proper time τ and weighted by $p_{\mathrm{LNV}} - \overline{p_{\mathrm{LNV}}}$, revealing the oscillatory pattern, for *two* HNLs with $M_N = 1\,\mathrm{GeV}$, $|\Theta_{\mu I}|^2 = 2 \cdot 10^{-8}$, $|\Theta_{eI}|^2 = |\Theta_{\tau I}|^2 = 0$ and $\delta M = 4 \cdot 10^{-7}\,\mathrm{eV}$.

in which case an accurate treatment of this uncertainty would be needed in the simulation. However, for longer oscillation periods like the one shown in figure 2.11, its effect should be negligible. Deriving precise sensitivity limits for HNL oscillations is beyond the scope of this paper, since it is likely that no simple analytical expression exists for them, due to the more complex test statistics required, compared to the detection or model-selection limits. HNL oscillations might for instance be amenable to methods such as maximum likelihood estimation, wavelets, or matched filtering, for which the null distribution can be estimated numerically using a (computationally expensive) bootstrapping procedure.

2.6 Conclusions

The SHiP experiment is set to have an unprecedented detection reach for a variety of models containing feebly interacting particles, such as Heavy Neutral Leptons (HNLs). A distinctive feature of SHiP among other intensity frontier experiments is its decay spectrometer, which allows it to not only place exclusion bounds, but also perform event reconstruction and measure the HNL properties. The simplest consistent HNL model accessible at SHiP contains two nearly degenerate HNLs, which can undergo oscillations. Their mass splitting δM is of particular interest, since it greatly influences their phenomenology as well as early-Universe cosmology (specifically, baryogenesis and dark matter production).

In the present work, we have investigated to which extent SHiP may be able to constrain or even measure δM. Depending on the scale of the oscillation phase $\delta M \tau$ accessible at an experiment, HNLs may or may not exhibit lepton number violation (LNV). The problem thus amounts to distinguishing LNC from LNV decay chains (figure 2.2) in a beam-dump setting (figure 2.1), where the primary lepton cannot be observed. We have shown that the angular distribution of the *visible* secondary decay products provides a partial solution to this problem, since, depending on the HNL mass, it can significantly differ between LNC and LNV in the laboratory frame (figure 2.7). This result has been qualitatively understood in the simplified case of two-body decays in the massless limit (figures 2.3 and 2.6). In order to handle more realistic cases, a Monte-Carlo simulation has been employed to generate accurate data sets of LNC and LNV events, including spin correlations and geometrical acceptance. The different distributions of the kinematic variables thus allow discriminating between LNC and LNV events using multivariate analysis; and with sufficiently many events, it becomes possible to statistically detect or exclude lepton number violation.

In order to produce sufficiently accurate training sets, our simulation must satisfy several requirements. It should be able to generate all the relevant two- and three-body meson decays containing an HNL (figure 2.4), as well as the selected HNL decay channel $N \to \pi^{\mp} l_{\beta}^{\pm}$. It should be accurate for GeV-scale HNLs, and should account for the spin correlations between the primary and secondary decays. Finally, it should run sufficiently fast to allow producing large training sets for various hypotheses and parameters. In order to meet all these requirements, we have written our own Monte-Carlo simulation, the output of which is used to train a binary classifier.

Knowing the accuracy of the classifier decision (figure 2.8) for a given mass and (relative) mixing angles, we can finally draw a "model-selection" sensitivity limit in the $(M_N, |\Theta|^2)$ plane (shown in figures 2.10a and 2.10b), above which SHiP should be able to either discover or rule out lepton number violation from HNLs. Interestingly, this limit lies below the detection sensitivity of near-future experiments such as NA62^{++}. This leads to a striking conclusion: SHiP might be able to not only *discover* HNLs, but also *characterize* them as either "Dirac-like" or "Majorana-like" fermions (depending on whether they feature LNV) even if previous experiments *see no signal at all.* Better yet, if the mass splitting between the two HNLs is of order $\delta M \sim 10^{-6}$ eV, SHiP should be able to *resolve the oscillations* of HNLs (figure 2.11), given sufficiently many events. Intriguingly, this mass splitting falls within the range required for producing dark matter in the νMSM [82]. Its measurement—or constraining—would therefore be an important test of cosmological models.

Acknowledgements

The authors would like to thank Mikhail Shaposhnikov for his suggestion of investigating HNL oscillations at SHiP, Oleg Ruchayskiy and Maksym Ovchynnikov for their helpful comments and proofreading of the present manuscript, Sonia Bouchiba and Federico Leo Redi for stimulating discussions about helicity effects in HNL decays, and Kyrylo Bondarenko and Elena Graverini for sharing with us their code to compute HNL decay widths. We thank the SHiP collaboration for their interest in our work. We are grateful to Annarita Buonaura, Richard Jacobsson, and Nicola Serra for their helpful comments and clarifications about the dimensions of the facility. JLT acknowledges support from the Carlsberg foundation. The work of IT was supported in part by the ERC-AdG-2015 grant 694896.

Appendix

2.A Simulation

2.A.1 Overview

It is not obvious whether the different angular correlations of LNC and LNV events lead to an observable effect in a realistic beam-dump experiment. To answer this question, we have devised a toy Monte-Carlo simulation, inspired from the one used in ref. [39], to simulate the production and decay of HNLs at the SHiP experiment [105, 106] (represented on figure 2.1).

The simulation of rare BSM processes with spin correlations entails two main requirements. First, we cannot afford to simulate all the possible processes, since, due to the small HNL mixing angles, the decay chains mediated by an HNL only represent a tiny fraction of all decays. Instead, we only simulate the BSM processes, and use importance sampling (i.e. introduce weights) in order to obtain the correct absolute number of events and expectation values (section 2.A.2).

Secondly, we cannot sample the primary and secondary decays separately, since they are not independent. Instead, we construct all possible decay chains for the production and decay processes of interest, and sample the entire chain at once, with a probability proportional to its combined branching fraction. The momenta of all the decay products are then sampled simultaneously, using the matrix element for the entire chain (section 2.A.4).

In addition, in order to accurately model the SHiP experiment, we need to sample the heavy meson momenta from a realistic spectrum (section 2.A.3) and take into account the finite size of SHiP and its geometrical acceptance (section 2.A.5). Finally, since most machine learning algorithms take unweighted data points as input, it is necessary to perform a last step of rejection sampling in order to produce a training set consisting of events with equal weights (section 2.A.6).

2.A.2 Decay chains

As discussed in ref. [129], the dominant HNL production process at SHiP is from weak decays of the lightest charmed or beauty mesons. In the present study, we focus on HNL masses below the D_s mass, and only select the fully reconstructible secondary decays $N \to \pi^\pm l_\beta^\mp$, By producing long-lived, charged particles which can be measured by the decay spectrometer located at the end of the decay vessel, they allow the HNL momentum to be reconstructed. The efficiency of particle identification at SHiP is high enough [130] that we can approximate it as one for the present estimate. Therefore we do not need to simulate decay chains containing any other secondary decays.

For the mixing angles of interest (i.e. below existing bounds), the fraction of all decays which are mediated by an HNL is tiny. We therefore need to use importance sampling in order to efficiently simulate only the processes of interest. For every proton on target (POT), the probability of producing a charmed hadron of species H is:

$$P(H) = \frac{\sigma_{cc}}{\sigma_{pN}} \cdot A_H \tag{2.19}$$

where σ_{cc} is the production cross-section for charmed hadrons, σ_{pN} the interaction cross-section for protons hitting the target nuclei, and A_H is the relative abundance of the charmed hadron species H (as given in appendix A of [106]). The *nominal* (i.e. physical) probability of producing an HNL which mediates a given decay chain $H \to [h']l_\alpha(N \to l_\beta h'')$ (irrespective of whether the decay is observed in the detector) is then:

$$\begin{aligned} P\left(H \to [h']l_\alpha(N \to l_\beta h'')\right) &= P(H) \cdot P(h' l_\alpha N | H) \cdot P(l_\beta h'' | h' l_\alpha N) \\ &= \frac{\sigma_{cc}}{\sigma_{pN}} \cdot A_H \cdot \mathrm{Br}_{\mathrm{prod}}(H \to [h']l_\alpha N) \cdot \mathrm{Br}_{\mathrm{decay}}(N \to l_\beta h'') \end{aligned} \tag{2.20}$$

where the last two terms are the production and decay branching ratios for HNLs in the considered decay chain. The *importance* distribution P' is defined as a uniform scaling for decay chains involving an HNL, and as zero for all other outcomes:

$$\begin{cases} P'\left(H \to [h']l_\alpha(N \to l_\beta h'')\right) = \frac{1}{w_{\mathrm{prod}}} P\left(H \to [h']l_\alpha(N \to l_\beta h'')\right) \\ P'(\text{no HNL}) = 0 \end{cases} \tag{2.21}$$

where w_{prod} is the weight to be applied to all the chains sampled from the importance distribution, and corresponds to the total probability of producing an HNL according to the nominal distribution:

$$w_{\mathrm{prod}} = \sum_{\text{chains}} P\left(H \to [h']l_\alpha(N \to l_\beta h'')\right) \tag{2.22}$$

When computing expected numbers of events over the entire duration of the SHiP experiment, which represents an integrated $N_{\mathrm{POT}} = 2 \cdot 10^{20}$ protons on target for 5 years of nominal operation, we must further multiply by N_{POT} the expectation values obtained for one event. This is most easily done by simply multiplying the total weights by N_{POT}.

2.A.3 Heavy meson spectrum

Once a chain is selected, we sample the momentum of the corresponding charmed meson from the spectrum measured by the LEBC-EHS collaboration [127] at the CERN SPS running at 400 GeV with a hydrogen target. The differential cross-section is parametrized as the product of a β distribution in x_F and an exponential distribution in p_T^2:

$$\frac{\mathrm{d}^2\sigma}{\mathrm{d}x_F \mathrm{d}p_T^2} = \sigma \frac{(n+1)b}{2} (1 - |x_F|)^n e^{-bp_T^2} \tag{2.23}$$

with the best-fit values $n = 4.9 \pm 0.5$ and $b = (1.0 \pm 0.1)\ \mathrm{GeV}^{-2}$. We thus implicitly assume the spectrum to be separable. Due to their very similar mass, and to compensate for the lack of data, we assume D_s mesons to share the same spectrum as D mesons.

By using the spectrum for a hydrogen target, we effectively neglect cascade production of heavy hadrons inside the target, leading us to underestimate the number of hadrons produced at the low-energy end of the spectrum. This could be problematic if their p_T spectrum happens to be significantly different from

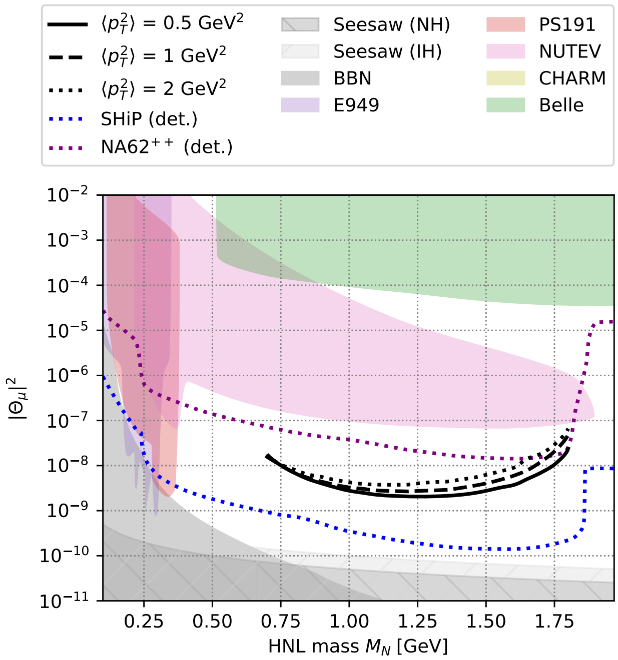

Figure 2.A.1: Effect of varying the width of the heavy meson p_T spectrum on the sensitivity to lepton number violation (90% CL), for an HNL coupling to the muon. Black lines represent the model-selection sensitivity of SHiP for various values of $\langle p_T^2 \rangle$. The dashed line corresponds to the best-fit value $\langle p_T^2 \rangle = 1\,\mathrm{GeV}^2$ from the LEBC-EHS collaboration [127].

that of primary hadrons produced in pp collisions. However, the lower acceptance for these softer hadrons should help mitigate the issue. In figure 2.A.1, we show how varying the width of the heavy meson p_T spectrum affects the final sensitivity. As expected, a larger p_T spread reduces the sensitivity, while a narrower spectrum improves it.

2.A.4 Decay product momenta

In order to preserve spin correlations between the HNL siblings and its decay products, we simulate both the HNL production and decay processes at once. For the masses and mixing angles of interest, the HNL is long-lived and can be assumed to be on its mass shell. Therefore the phase-space sampling can be performed independently for the primary and secondary decays. We use the m-generator algorithm [136] for that, as described in ref. [137]. In order to sample events with a probability proportional to the squared transition amplitude, we then perform rejection sampling, taking the phase-space distribution as proposal distribution, and an acceptance probability proportional to the spin-summed, squared matrix elements (2.12) and (2.13) for the entire decay chain. Only the spin states of the external particles (which interact with the detector and are thus "measured" in the quantum mechanical sense) are summed over.

2.A.5 Geometry

In order to model the geometry of the SHiP experiment, we must account for the finite size of the detector and its geometrical acceptance. In the current SHiP design (represented on figure 2.1), the fiducial volume consists of an evacuated right pyramidal frustum of length 50 m, located at a distance of 50 m from the target, and with horizontal and vertical sides 5 m and 10 m respectively at the far end. It is followed by a 10 m long tracking station.

To estimate the probability of the HNL decaying within the fiducial volume and passing the acceptance cuts, we use once again importance sampling for sampling the decay vertex. This is required in order to overcome the potentially very long lifetime of HNLs, which could cause most of them to decay away from the experiment. We choose an importance distribution (approximately) covering the fiducial volume, by sampling the decay vertex uniformly along the HNL momentum, at a distance such that it falls inside the decay vessel. The nominal decay probability density is, as a function of the proper time τ (or boost factor γ and distance L) between the HNL production and decay:

$$P_{\text{decay}}(\tau) = \Gamma e^{-\Gamma\tau} \quad \Longrightarrow \quad P_{\text{decay}}(L|\gamma) = \frac{\Gamma}{\beta\gamma} e^{-\frac{\Gamma L}{\beta\gamma}} \tag{2.24}$$

The partial weight resulting from this importance sampling step is therefore:

$$w_{\text{decay}}(L|\gamma) = \frac{\Gamma L_{\text{DV}}}{\beta\gamma\cos(\theta)} e^{-\frac{\Gamma L}{\beta\gamma}} \tag{2.25}$$

where $L_{\text{DV}} = 50\,\text{m}$ is the length of the decay vessel and θ the angle between the HNL momentum and the beam axis. In the linear regime, where $\Gamma\tau \ll 1$, this partial weight reduces to $w_{\text{decay}}(L|\gamma) \cong {}^{\Gamma L_{\text{DV}}}/_{\beta\gamma\cos(\theta)}$.

We finally apply acceptance cuts by requiring the HNL to decay within the decay vessel, and the trajectories of its two decay products ($l_\beta^\mp$ and $\pi^\pm$) to intersect the tracking station located at its far end.

2.A.6 Unweighting

As a last step, we perform again rejection sampling on the weighted events in order to obtain a set of events with equal weights, which are easier to analyse and process with machine learning algorithms. This is done by accepting events with a probability proportional to their weight, and can be justified as follows.

Let X denote a random variable representing the simulated event, and x a concrete realization of it. Let $f(x) = P(X = x)$ be the nominal (i.e. true) distribution and $g(x)$ the importance distribution, such that $g(x) > 0$ for all outcomes x in the domain of interest Ω (i.e. all relevant observables must have their support in Ω). If x is sampled from the importance distribution $g(x)$, its associated weight will be $w(x) = f(x)/g(x)$. Let M be an upper bound on $w(x)$, i.e. $M \geq w(x), \forall x \in \Omega$. If we choose the acceptance probability to be $a(x) \overset{\text{def}}{=} w(x)/M \leq 1$, then it immediately follows that the accepted events, effectively drawn from the new importance distribution $g(x) \cdot a(x)$, will have uniform weight M.

It is therefore possible to perform rejection sampling a posteriori in order to produce uniformly weighted events. However, storing all the generated events, many of which will eventually be rejected, would be inefficient from a memory perspective. A more economical solution, which we decided to use, consists in performing rejection sampling directly as events are being generated. This requires estimating an upper bound M on the weights, during an initial burn-in phase.

2.B LNC/LNV classification

At leading order in the light lepton and hadron masses, the matrix elements for LNC and LNV decay chains have a straightforward analytical dependence on the invariant mass s_{ll} of the charged lepton pair. However, unlike in collider experiments, this variable is not readily available in a beam-dump setting, due to the primary lepton being unobservable. As we saw in section 2.3.2, the different angular correlations between the charged leptons can nevertheless lead to residual correlations between the visible HNL decay products. The absence of an obvious test statistics, along with the almost background-free conditions and highly efficient PID at SHiP [130], makes the task of distinguishing LNC from LNV ideally suited for multivariate analysis. In the following subsections, we describe how we generate the training set (section 2.B.1), the classifier used to discriminate between LNC and LNV events (section 2.B.2), how to produce a sensitivity limit from its output (section 2.B.3), and finally how sensitive is the classification to systematic uncertainties on the heavy meson spectrum (section 2.B.4).

2.B.1 Dataset

As mentioned in section 2.4.2, we need to generate datasets for various HNL masses M_N and rays in $|\Theta_\alpha|^2$ space, where $\alpha = e, \mu$ (the overall normalization does not matter). In practice, we choose a mass range spanning the region between the K and D_s thresholds, and consider several benchmark models with fixed $|\Theta_e|^2 : |\Theta_\mu|^2$ ratios.[11] For each choice of physical parameters, we sample $3 \cdot 10^6$ events with uniform weights. This is done by sampling sufficiently many weighted events and, as they are being generated, "unweighting" them by performing rejection sampling with an acceptance probability proportional to their weight. Only events which pass the acceptance cuts are used for training. In the simulation, the HNL is taken to be a single Majorana particle, such that the dataset contains equal numbers of LNC and LNV events and is balanced with respect to the primary and secondary lepton charges. We select only the fully reconstructible HNL decays $N \to \pi^\mp l_\beta^\pm$, which do not contain an unobservable light neutrino, and produce long-lived charged particles which can be measured by the decay spectrometer. For the sake of simplicity, we will assume the PID to be perfectly efficient throughout this analysis. Non-trivial efficiencies are expected to slightly reduce the final sensitivity reach. As explained in section 2.4.2, each event is labelled as being either LNC or LNV,

[11] We do not consider HNL production through τ mixing in this work, since it would have required to implement secondary production from τ decays. It is negligible in the considered mass range unless the Θ_τ mixing angle is significantly larger than the others, as can be seen in figure 2.4. In addition, visible HNL decays through τ mixing are forbidden below the τ threshold.

Feature	p2y	p3y	p2x	p3x	pCMz	zD	xD	yD	p1x	pCMy
# splits	302	282	243	238	141	114	105	97	91	85

Feature	pCMx	p1y	E1	E2	E3	p3z	p2z	p1z	Q12
# splits	77	74	69	67	61	53	34	14	9

Table 2.B.1: Feature importance for a 1 GeV HNL coupling to μ.

and we record the 19 observable features listed in table 2.1. The dataset is split into training / validation / test sets with respective proportions $0.5 : 0.2 : 0.3$.

2.B.2 Classifier

We employ the LightGBM [131] gradient boosting algorithm, accessed through the Python interface to the reference implementation [132]. For classification, we choose the `binary` objective. We use early stopping based on the binary log-loss (`binary_logloss`) and the area-under-curve (`auc`) metrics, with a 10 round threshold. The hyperparameters `num_leaves` and `learning_rate` are manually optimized by maximizing the above two metrics on the validation set. The classification accuracy is presented in figure 2.8 as a function of the HNL mass M_N for two orthogonal scenarios, corresponding to the HNL coupling exclusively to electrons ($|\Theta_e|^2 : |\Theta_\mu|^2 : |\Theta_\tau|^2 = 1 : 0 : 0$) or muons ($|\Theta_e|^2 : |\Theta_\mu|^2 : |\Theta_\tau|^2 = 0 : 1 : 0$), and a third one where it couples equally to both ($|\Theta_e|^2 : |\Theta_\mu|^2 : |\Theta_\tau|^2 = 1 : 1 : 0$).

It is instructive to understand the origin of this dependence, if only to make sure that it corresponds to a physical effect. LightGBM provides a way to estimate the feature importance, by counting the number of times a feature is used to split a tree. Those are listed in table 2.B.1 for a 1 GeV HNL coupling to muons (which results in a classification accuracy of 63.5%). They reveal that the most important features are the transverse components of the momenta of the HNL decay products. Indeed, it is possible to successfully train a model using a single feature such as the transverse momentum $p_{T,\mu}$ of the secondary muon, while still obtaining a classification accuracy of 61.5% (for the same dataset).

Inspecting the results more closely (see figure 2.7) shows that LNV events have on average a slightly larger transverse momentum than LNC ones. This is consistent with our discussion from section 2.3.2, and allows us to understand the mass dependence. At large HNL masses, as we approach the closing mass of D meson leptonic decays, the kinetic energy of the HNL in the heavy meson rest frame decreases, until it becomes so small that the difference between LNC and LNV becomes negligible compared to the transverse momentum spread of the heavy meson spectrum. As the HNL mass decreases, 3-body semileptonic decay channels open, and become dominant at lower masses. The additional meson takes away part of the energy from the HNL, leaving it with insufficient kinetic energy to “escape” the transverse momentum spread of the heavy meson spectrum. Finally, the large boost of the heavy mesons along the beam axis washes out most of the information contained in the longitudinal part of all laboratory frame momenta, which explains their low importance.

2.B.3 Sensitivity to lepton number violation

As stated in section 2.4.3, our main goal is to distinguish between the following two hypotheses using exclusively the classifier decision (i.e. not the underlying feature vector z):

- $\mathcal{H}_1$: HNLs are Dirac or quasi-Dirac with $\delta M\tau \ll 1$ (LNC decays only).
- $\mathcal{H}_2$: HNLs are Majorana or quasi-Dirac with $\delta M\tau \gg 1$ (LNC and LNV decays).

Those can be expressed as special cases of a more general hypothesis $\mathcal{H}(f)$, $f \in [0,1]$, parametrized by the relative frequency f of LNV events:

- $\mathcal{H}(f)$: (LNV rate) $= f \times$ (total rate).

such that $\mathcal{H}_1 = \mathcal{H}(f=0)$ and $\mathcal{H}_2 = \mathcal{H}(f=1/2)$.

We model the classifier decisions using a 2×2 confusion matrix $C_{ij} = P(i \text{ classified as } j)$, where $i,j = 1,2$ correspond to the two classes, respectively LNC and LNV. The confusion matrix can be expressed in terms of the classification accuracies as:

$$C = \begin{pmatrix} a_1 & 1-a_1 \\ 1-a_2 & a_2 \end{pmatrix} \tag{2.26}$$

Suppose we observe N events passing the selection cuts, k of which are classified as LNV. Then, under $\mathcal{H}(f)$, the likelihood of classifying $N-k$ events in class 1 (LNC) and k in class 2 (LNV) is given by the following binomial distribution:

$$\mathcal{L}(k;f) = \binom{N}{k}\big(a_2 f + (1-a_1)(1-f)\big)^k \big(a_1(1-f) + (1-a_2)f\big)^{N-k} \tag{2.27}$$

Under hypothesis $\mathcal{H}_1$, i.e. all events are LNC, this likelihood reduces to:

$$\mathcal{L}_1(k) = \mathcal{L}(k; f=0) = \binom{N}{k}(1-a_1)^k a_1^{N-k} \tag{2.28}$$

while under hypothesis $\mathcal{H}_2$, i.e. events come from either class with equal probability, it becomes:

$$\mathcal{L}_2(k) = \mathcal{L}(k; f=1/2) = \binom{N}{k}\frac{(1+a_2-a_1)^k(1+a_1-a_2)^{N-k}}{2^N} \tag{2.29}$$

For many models, including LightGBM (with a balanced training set), $a_1 \approx a_2 \stackrel{\text{def}}{=} a$. In this limit, $\mathcal{L}_2(k)$ simplifies to $\binom{N}{k}2^{-N}$.

Since $\mathcal{H}_{1,2}$ and $\mathcal{H}(f)$ are nested, then, assuming we have sufficiently many events, we can use Wilk's theorem[12] to try to exclude $\mathcal{H}_{1,2}$. To this end, we construct the two likelihood ratios $\Lambda_{1,2}(k)$ as:

$$\Lambda_i(k) = \frac{\mathcal{L}_i(k)}{\mathcal{L}(k;\hat{f})}, \quad i = 1,2 \tag{2.30}$$

[12] A potential issue in the case of $\mathcal{H}_1$ could be that the null value $f=0$ lies on the boundary of the domain $[0,1]$ of f, while Wilk's theorem requires the true value to be in the interior of the parameter space. However, $\ln(\mathcal{L}(k;f))$ has a well-behaved analytical continuation over a domain larger than $[0,1]$. As long as the estimator $\hat{f}$ has a sufficiently small variance, this boundary effect can therefore be ignored and Wilk's theorem still applies. See [138] for a comprehensive discussion of the validity conditions of Wilk's theorem.

where $\hat{f}$ is the maximum likelihood estimator for f:

$$\hat{f} = \frac{1 - a - k/N}{1 - 2a} \tag{2.31}$$

Wilk's theorem states that if $\mathcal{H}_i$ ($i = 1$ or 2) is realized, then $-2\ln(\Lambda_i(k))$ follows a χ^2 distribution with one degree of freedom. Conversely, if we observe $-2\ln(\Lambda_i(k)) > 2.7$, then $\mathcal{H}_i$ will be disfavoured at 90% CL. If both hypotheses $\mathcal{H}_{1,2}$ were disfavoured simultaneously, this would suggest $\delta M\tau \sim 2\pi$ and potentially resolvable HNL oscillations.

If hypothesis $\mathcal{H}_1$ is actually realized, we expect k to take a value around the expected number of events misclassified as LNV: $(1-a)N$, which, for large N, is approximately equal to the median. The median of the log-likelihood-ratio when testing for $\mathcal{H}_2$ is therefore:

$$\mathrm{med}_1\left(\ln(\Lambda_2)\right) \approx -N \underbrace{\left(\ln(2) + a\ln(a) + (1-a)\ln(1-a)\right)}_{\stackrel{\text{def}}{\equiv}\, l_1(a)>0} \tag{2.32}$$

If, instead, $\mathcal{H}_2$ is realized, then we expect k to take a median value of approximately $N/2$, such that:

$$\mathrm{med}_2\left(\ln(\Lambda_1)\right) \approx N \underbrace{\left(\ln(2) + \frac{1}{2}\ln(a) + \frac{1}{2}\ln(1-a)\right)}_{\stackrel{\text{def}}{\equiv}\, l_2(a)<0} \tag{2.33}$$

For a fixed confidence level, we can invert these two formulas to estimate, for each true hypothesis $\mathcal{H}_i$, $i = 1, 2$, the median number of events $N_i(a)$ required to exclude the other hypothesis:

$$N_i(a) = \left|\frac{\ln(\Lambda_{\mathrm{cr}})}{l_i(a)}\right| \tag{2.34}$$

with $-2\ln(\Lambda_{\mathrm{cr}}) \approx 2.7$ for a 90% CL. The higher the classification accuracy, the less events are required to reach the target, while accuracies close to 1/2 do not allow distinguishing the two hypotheses, as $N_i(1/2) \to \infty$. So far we have only considered the two extreme cases $f = 0$ or $1/2$, i.e. $\delta M\tau \lesseqgtr 2\pi$. We can generalize this analysis to the case where the true hypothesis or the null hypothesis have a non-trivial LNV fraction f. A larger number of events will then be required to reach the same confidence level. We will not discuss these cases further in this paper, in order to avoid making the discussion unnecessarily complicated.

As a final step, for each HNL mass M and ratio $\left|\Theta_e\right|^2 : \left|\Theta_\mu\right|^2 : \left|\Theta_\tau\right|^2$, we compute the squared mixing angles $\left|\Theta_\alpha\right|_i^2(M)$ required to produce $N_i(a(M))$ events, thus producing for each true hypothesis $\mathcal{H}_i$ a sensitivity limit, above which SHiP should be able to exclude the other hypothesis with a probability of at least 1/2. The resulting sensitivity plots are presented in section 2.5.1.

2.B.4 Systematic uncertainties coming from the heavy meson spectrum

For a classifier to generalize well out of sample, i.e. on real-world data, the distribution used for training should match the true, physical distribution of features.

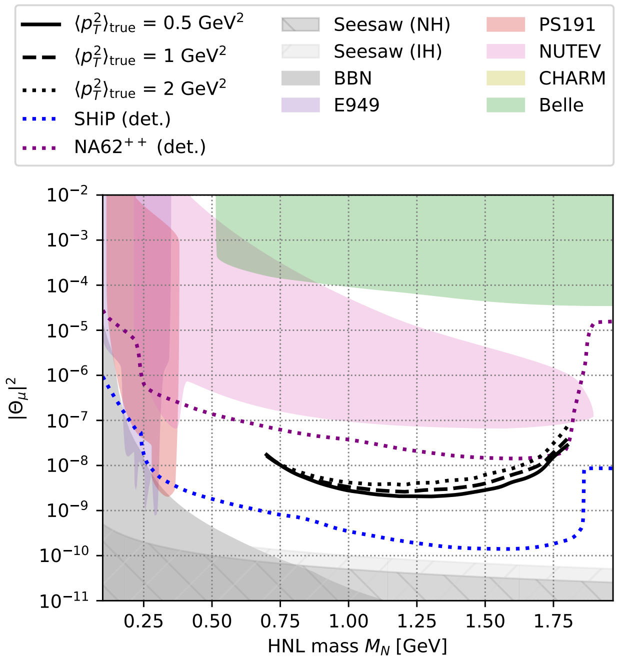

Figure 2.B.1: Effect on the LNV sensitivity (90% CL) of computing the classification accuracy on a test set generated with a different p_T spectrum compared to the training set, for an HNL coupling to the muon. Black lines represent the model-selection sensitivity of SHiP for various true $\langle p_T^2 \rangle$. Here, the training set is always generated with $\langle p_T^2 \rangle = 1\,\text{GeV}$.

This is in general not the case, since a simulation never perfectly represents reality. We can, however, work around this requirement by explicitly evaluating the classification accuracy over a set of test distributions which is likely to encompass the true distribution. This requires knowing and parametrizing the uncertainties coming from the simulation. We can then obtain a conservative estimate for the classification accuracy by varying the unknown parameters within their uncertainties, and taking a lower bound. If this lower bound is high enough, we should still be able to probe lepton number violation on real data.

At SHiP, the main uncertainty affecting the LNC / LNV classification accuracy comes from the transverse momentum spread of the heavy meson spectrum, which is only known with limited accuracy. In order to estimate the actual sensitivity of SHiP to LNV for a realistic dataset, we therefore compute the classification accuracy for a family of test sets generated using slightly different p_T spectra, and we take the lowest value as our estimate. The change in the sensitivity resulting from varying $\langle p_T^2 \rangle$ by a factor of two up and down with respect to the best-fit value from LEBC-EHS [127] is shown in figure 2.B.1. The planned charm spectrum measurements at SHiP should be able to constrain $\langle p_T^2 \rangle$ to a much better accuracy than the range displayed in the figure.

Interestingly, when comparing this result with figure 2.A.1, we observe that the classification accuracy seems to mostly depend on the $\langle p_T^2 \rangle$ of the test set, but not much on the one used for training. This suggests that we might be able to safely use the best-fit spectrum for training without worrying about biasing the results should the true spectrum turn out to be different, provided that we use a conservative estimate for the accuracy. In a more comprehensive study, one would likely want to vary additional parameters related to the spectrum, geometry and simulation.

Chapter 3

Projected NA62 sensitivity to heavy neutral lepton production in $K^+ \to \pi^0 e^+ N$ decays

Authors: Jean-Loup Tastet, Evgueni Goudzovski, Inar Timiryasov, Oleg Ruchayskiy

ArXiv: 2008.11654

Abstract

Heavy neutral leptons (HNLs) appear in many extensions of the Standard Model of particle physics. In this study, we investigate to which extent the NA62 experiment at CERN could improve the existing bounds on the HNL mixing angle $|U_e|^2$ by performing a missing mass search in $K^+ \to \pi^0 e^+ N$ decays in flight. We show that the limit $|U_e|^2 \simeq 2 \times 10^{-6}$ can be reached with the currently available data in the mass range $125 - 144\,\mathrm{MeV}$, which is currently not well covered by production searches. Future data, together with a dedicated trigger and/or improvements in rejection of out-of-acceptance photons, can improve this limit by another order of magnitude.

Contents

3.1 Introduction

3.1.1 Heavy Neutral Leptons

Despite its astounding success in describing the outcomes of collider experiments, the Standard Model of particle physics (SM) fails to account for multiple reliable observations: the baryon asymmetry of the Universe (BAU, see e.g. ref. [23]), dark matter (see e.g. ref. [139]), as well as neutrino flavor mixing and oscillations [140]. The latter observations provide unambiguous evidence for non-zero neutrino masses, which call for the introduction of additional degrees of freedom into the SM. Among many models explaining neutrino masses, those that introduce no new particles above the electroweak scale are of special interest, since they do not destabilize the Higgs mass [89, 90, 141] and are accessible already by the current generation of experiments (see e.g. ref. [91]).

Such particles may appear for example in extensions of the neutrino sector (see e.g. refs. [106, 142]) such as the type-I seesaw theories [27, 30, 31, 33, 34, 88]. The assignment of charges in the SM predicts that hypothetical right-handed counterparts to neutrinos would be completely neutral, i.e. transform as singlets under the SM gauge group. As such, they also admit a Majorana mass term whose value is not predicted from neutrino data. The physical spectrum of these theories contains three light neutrino mass states $\nu_{\mathrm{L}i}$ plus a number of new *heavy neutral leptons* (HNLs) $N_{\mathrm{R}I}$ (conventionally defined as right-handed to be consistent with other $SU(2)_L$ singlet fermions). These heavy neutral leptons inherit from the active neutrino flavor states their weak-like interactions with W and Z bosons, albeit with a coupling suppressed by the (flavour-dependent) elements of the mixing matrix $\Theta_{\alpha I} \ll 1$. In what follows, we will refer to the elements of this matrix as *mixing angles*. The active neutrino flavors $\nu_{\mathrm{L}\alpha}$ ($\alpha = e, \mu, \tau$) are then a superposition of light and heavy mass states: $\nu_{\mathrm{L}\alpha} = V_{\alpha i}^{\text{PMNS}} \nu_{\mathrm{L}i} + \Theta_{\alpha I} N_{\mathrm{R}I}^c$, where $V_{\alpha i}^{\text{PMNS}}$ is the (now non-unitary) PMNS matrix (see e.g. ref. [143]).

HNLs can by themselves resolve the aforementioned beyond-the-Standard-Model puzzles, as in the *Neutrino Minimal Standard Model* (νMSM) [35, 36]. Or they can serve as a *portal* (mediator) between the SM sector and other hypothetical sectors containing new particles (see e.g. refs. [91, 106] and references therein). In the latter case HNLs can possess other types of interactions (see e.g. refs. [31, 32, 114, 144–153]), in addition to those inherited from their mixing with the active flavor states.

In this paper, we consider a simplified model containing one HNL N with three flavour mixing angles $U_\alpha \ll 1$. It can be thought either as a single Majorana mass state, or several HNLs degenerate in mass, in which case the equivalent mixing angle that we constrain is $|U_\alpha|^2 = \sum_I |\Theta_{\alpha I}|^2$.

3.1.2 Missing mass searches

Intensity frontier experiments like NA62 at CERN are, thanks to the high statistics available, well suited to constrain HNLs. There are two main experimental methods to search for HNLs: production and decay searches [91]. Production searches consist in reconstructing the "missing" momentum of invisible particles from an otherwise known kinematical configuration, and searching for a mass peak emerging over a smooth background — which indicates the presence of a new particle. They can be performed only if the kinematics of the process are fully known, as e.g. at kaon factories or e^+e^- colliders. Decay searches consist in identifying visible final states in the HNL decays and can be performed at fixed-target, beam dump, e^+e^- or pp collider based experiments. Production searches are sensitive to the HNL production rate alone, but not to its lifetime[1] or decay modes. In typical models, the production rate is proportional to the square of a single mixing angle active in the production process, $|U_\alpha|^2$ ($\alpha = e$, μ or τ). A non-observation can therefore be directly translated into a limit on this mixing angle, with little model dependence. On the other hand, decay searches are sensitive to a combination of the various squared mixing angles involved in the HNL production, multiplied by the partial HNL decay width[2], which in typical models also depends on squared mixing angles. The signal is thus proportional to a combination of fourth powers of mixing angles. In non-minimal models, it will depend on additional parameters. To be translated into a set of exclusion limits, a non-observation must therefore be interpreted within a specific model to disentangle the contributions of the various flavors, hence introducing model dependence.

3.1.3 The NA62 experiment

The NA62 experiment at CERN [104] employs a high intensity, almost monochromatic secondary K^+ beam of 75 GeV momentum to measure the rate of the ultra-rare $K^+ \to \pi^+\nu\bar{\nu}$ decay to a 10% precision using the decay in flight technique [104]. The beam is delivered into a 80 m long vacuum tank, giving rise to a K^+ decay rate in the tank of about 5 MHz. Both the incoming kaons and their visible decay products are detected, allowing to reconstruct the missing momentum. The experiment is equipped with a system of veto detectors for both charged and neutral particles. In particular, the photon veto helps reducing the contribution of undetected photons and π^0 mesons to the missing momentum. This leads to favourable background conditions, and provides sensitivity to K^+ decays with invisible particles in the final state, which are reconstructed using the *missing mass* technique. Such searches have been performed [154, 155] or are planned both for HNLs and for other feebly interacting particles [42, 156–159].

The NA62 collaboration has recently performed a search for HNL (N) production in the $K^+ \to e^+N$ decay with the full Run 1 (2016 – 2018) data set, and established stringent limits at the level of $|U_e|^2 \sim 10^{-9}$ in the HNL mass range 144 – 462 MeV [155]. The sensitivity of this search deteriorates abruptly at lower HNL masses due to the shape of the background. On the other hand, the upper limits on $|U_e|^2$ established by the searches for the $\pi^+ \to e^+N$ process are at the 10^{-8} level up to a mass of about 120 MeV, and weaken sharply above this point [160]. As a result, *production* searches only weakly constrain the mass range 120 – 144 MeV, and they do not currently provide any constraints in the range 135 – 144 MeV.

[1] Provided it is long enough that the HNL does not decay within the experimental setup.
[2] Unless the HNL decays promptly, in which case it is the branching fraction that matters.

3.1.4 Previous bounds

It should be noted that beam dump experiments, notably PS191 at CERN [161], have obtained competitive bounds in the above mass range (see also ref. [162] for a re-analysis including the neutral current contribution). The PS191 experiment was designed specifically to detect decay products of heavy neutrinos in a low-energy neutrino beam produced by kaon and pion decays. Such a method hinges on the assumption that not only production, but also the visible decay of HNLs is determined by the mixing angles $|U_\alpha|^2$ and $|U_\beta|^2$, where α and β may be the same or different flavours (the so-called $|U|^4$ experiments, see e.g. refs. [106, 162]). The missing mass searches, although less sensitive in the case of the minimal (type I) seesaw model, are applicable to a wider class of models. In particular, those models where, due to other interactions, HNLs decay before reaching the PS191 detector (~ 70 meters away from the target) and therefore evade PS191 constraints, can still be probed by missing mass searches. This motivates the present study, which consists in probing the 120 – 144 MeV mass range for the electron mixing at NA62 using the missing mass technique in the $K^+ \to \pi^0 e^+ N$ channel.

3.2 Signal simulation

The proposed search involves the final state consisting of a positron and two photons originating from a prompt π^0 decay. The expected number of signal events is

$$s_{\text{tot}} = N_K \times \text{BR}(K^+ \to \pi^0 e^+ N) \times \epsilon_{\text{sig}} \times (1 - P_{\text{decay}}), \tag{3.1}$$

where N_K denotes the effective number of K^+ decaying within the fiducial volume, P_{decay} is the probability that the HNL decays *visibly* inside the detector (in which case the event is ignored by the present analysis), and ϵ_{sig} is the signal detection efficiency (including the geometrical acceptance, but not the probability of the HNL decaying outside the detector).

P_{decay} is a model-dependent parameter determined by the specific HNL decay channels. However, for sufficiently long-lived HNLs, $1 - P_{\text{decay}} \approx 1$ and this factor can be omitted. The matrix element of the decay, and the branching ratio $\text{BR}(K^+ \to \pi^0 e^+ N)$, both of which depend on the assumed HNL mass m_N, are computed following refs. [129, 163], using the measured form factors from ref. [164]. The branching ratio is shown as a function of m_N as the blue dashed line in figure 3.2.1. The NA62 Run 1 data sample currently available for the $K^+ \to \pi^0 e^+ N$ search, collected using a 1-track trigger with an effective prescaling factor of about 150, corresponds to $N_K \approx 3 \times 10^{10}$ [155]. The acceptance ϵ_{sig} is computed by interfacing our matrix element sampler with the full Geant4-based NA62 simulation framework [165], and employing a basic event selection requiring a positron and two photons from a $\pi^0 \to \gamma\gamma$ decay in the geometric acceptance of the detector. The acceptance is found to be about 10% for HNL masses below 150 MeV, and to decrease as a function of m_N for higher masses towards the kinematic endpoint.

The events are binned in squared missing mass $m^2_{\text{miss}} = (p_{K^+} - p_{\pi^0} - p_{e^+})^2$. The finite momentum and energy resolution of the detector causes the reconstructed signal m^2_{miss} distribution to follow a Gaussian profile centered at

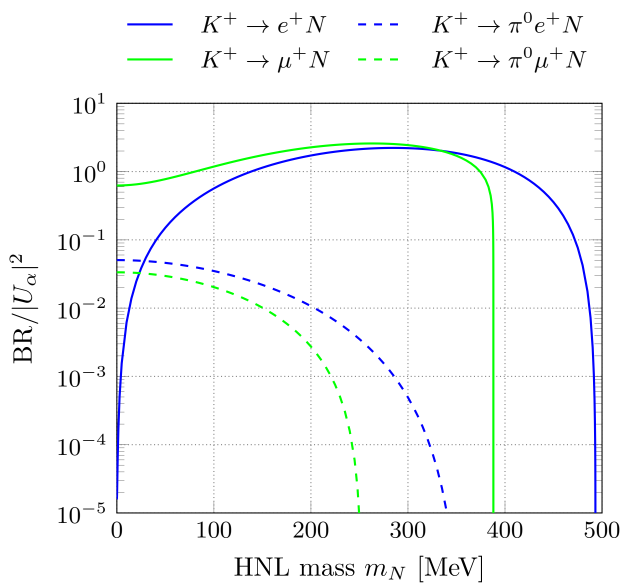

Figure 3.2.1: HNL production branching ratios in leptonic and semileptonic K^+ decays, normalised to the squares of the relevant mixing angles. Contrary to the $K^+ \to e^+N$ decay, the $K^+ \to \pi^0 e^+ N$ decay considered in this study is not helicity-suppressed for $m_N \ll m_{K^+}$.

m_N^2. The typical NA62 resolution on m^2_{miss} is about $10^{-3}\,\text{GeV}^2$. A value of $\sigma_{m^2} = 1.7 \cdot 10^{-3}\,\text{GeV}^2$ obtained for the $K^+ \to e^+\nu$ decay [155] is assumed conservatively for this study.

3.3 Background estimate

The dominant source of background to the $K^+ \to \pi^0 e^+ N$ process comes from the radiative $K^+ \to \pi^0 e^+ \nu_e \gamma$ inner-bremsstrahlung decay with the radiative photon escaping detection, thus causing m^2_{miss} to be mis-reconstructed. The expected reconstructed m^2_{miss} spectrum of the $K^+ \to \pi^0 e^+ \nu_e \gamma$ process, simulated according to ref. [166], taking into account the NA62 acceptance and resolution, and assuming that the radiative photon is not detected, is shown in figure 3.3.1. The principal contribution to the background for HNL masses above 100 MeV comes from the radiative tail, while the contributions from the main $K^+ \to \pi^0 e^+ \nu_e$ peak at $m^2_{\text{miss}} = 0$ (caused by the finite mass-squared resolution) and non-Gaussian reconstruction tails are subleading. The origin and properties of this background are similar to those encountered in the search for the $K^+ \to \mu^+ N$ decay at NA62 [167].

Other background sources, such as $K^+ \to \pi^0\pi^0 e^+ \nu_e$ decays with both photons from a π^0 decay evading detection, or $K^+ \to \pi^0 \mu^+ \nu_\mu$ decays followed by $\mu^+ \to e^+ \bar{\nu}_\mu \nu_e$ decays, are found to be subleading. In particular, misreconstruction of the $K^+ \to e^+$ decay vertex position in the latter case typically leads to the invariant mass of the two photons from the $\pi^0 \to \gamma\gamma$ decay, reconstructed assuming photon emission at the decay vertex, being incompatible with the π^0

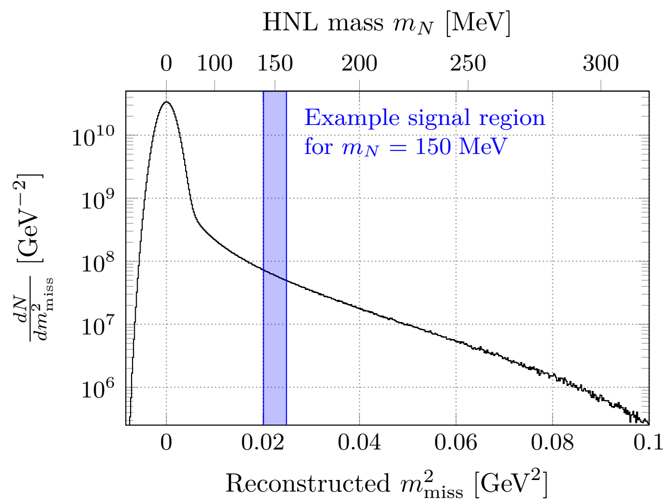

Figure 3.3.1: Reconstructed squared missing mass spectrum of the $K^+ \to \pi^0 e^+ \nu_e \gamma$ background, obtained by modelling the NA62 acceptance and resolution, and assuming that the radiative photon is not detected. The total number of reconstructed events in the spectrum is 1.5×10^8, corresponding to $N_K = 3 \times 10^{10}$ kaon decays considered. A $\pm 1.4\sigma_{m^2}$ wide signal region for $m_N = 150$ MeV is shown for illustration.

mass.

The background from radiative photons is largely reducible thanks to the NA62 photon veto system, which provides hermetic geometric coverage for photon emission angles θ_γ up to $\theta_{\mathrm{max}} = 50$ mrad with respect the beam axis, and partial geometric coverage (of approximately $\theta_{\mathrm{max}}/\theta_\gamma$) for larger emission angles. The nominal detection inefficiency for energetic photons is 10^{-3} for the large-angle system (for photon energies in excess of a few hundred MeV), and well below 10^{-3} for the intermediate and small angles [104]. As seen in figure 3.3.2, most of the photons from $K^+ \to \pi^0 e^+ \nu_e \gamma$ decays susceptible to contaminate the relevant signal regions (for $m_N \gtrsim 100$ MeV, i.e. $m^2_{\mathrm{miss}} \gtrsim 0.01$ GeV^2) are emitted within 50 mrad of the beam axis. A simplified photon detection efficiency model is used in this study: the nominal detection inefficiency of 10^{-3} is assumed for $\theta_\gamma < \theta_{\mathrm{max}}$ (this assumption is valid as the energy of the photons intercepting the large-angle veto acceptance for the m^2_{miss} range of interest is always above 200 MeV, and is typically in the GeV range), and zero detection efficiency is assumed conservatively for the (softer) photons emitted at $\theta_\gamma \geq \theta_{\mathrm{max}}$. In this model, the background events are dominated by those with soft photons emitted at angles above 50 mrad outside the hermetic coverage zone. Therefore the accuracy of the detection efficiency model does not significantly affect the background estimate.

3.4 Projected NA62 sensitivity

To estimate the projected sensitivity, we use for simplicity a cut-and-count analysis. We expect that the actual search will instead involve spectrum shape

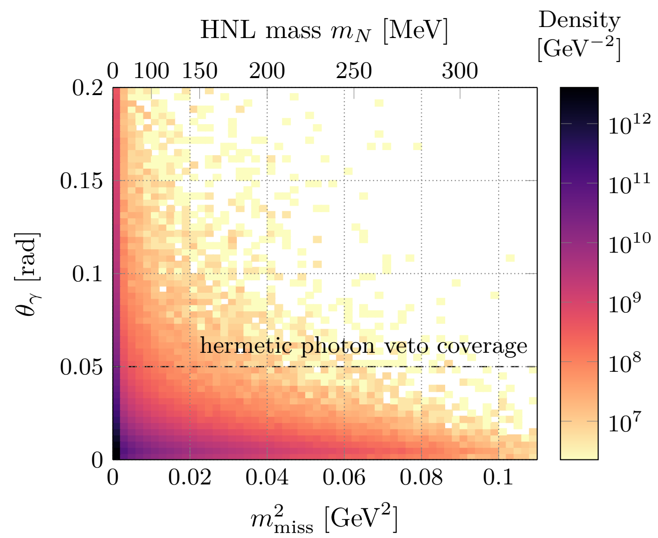

Figure 3.3.2: $K^+ \to \pi^0 e^+ \nu_e \gamma$ background event density as a function of the true missing mass squared and angle θ_γ between the photon and the beam axis. Hermetic geometric coverage is provided for photons with $\theta_\gamma < 50\,\text{mrad}$.

analysis. We define the signal region for a HNL of mass m_N as a rolling window of missing mass squared $m^2_{\text{miss}} \in [m^2_N - k\sigma_{m^2}, m^2_N + k\sigma_{m^2}]$, where the width $\sigma_{m^2} = 1.7 \cdot 10^{-3}\,\text{GeV}^2$ corresponds to the approximate mass-squared resolution of the detector [155] and the constant $k = 1.4$ is chosen to maximize the $s/\sqrt{b}$ ratio (where s and b respectively represent the numbers of signal and background events) and therefore the power of the search. A typical signal region is shown in figure 3.3.1. Real photon emissions produce a smoothly falling background in m^2_{miss}. The search is performed by looking for a significant excess of events over the background count b in each signal region. The detection sensitivity is expressed as a 90% confidence limit (local significance), which roughly corresponds to $s \gtrsim 1.282\sqrt{b}$ in the limit $b \gg 1$, with $s = \text{erf}(k/\sqrt{2}) \times s_{\text{tot}}$ the approximate number of signal events inside the signal region. The projected, median *exclusion* limit is similarly obtained, by replacing $\sqrt{b}$ with $\sqrt{b+s}$. The background b from real photon emissions, integrated over a small m^2_{miss} window, is approximately:

$$b(m^2_{\text{miss}}) \approx 2k \times \sigma_{m^2} \times \langle \epsilon_{\text{bkg}} \rangle \times \frac{\mathrm{d}N(K^+ \to \pi^0 e^+ \nu_e \gamma)}{\mathrm{d}m^2_{\text{miss}}} \tag{3.2}$$

where $\langle \epsilon_{\text{bkg}} \rangle$ denotes the mean background efficiency of the veto system in this window. This results in a *detection* sensitivity of:

$$|U_e|^2 \gtrsim \frac{2.56}{N_K} \frac{\sqrt{\langle \epsilon_{\text{bkg}} \rangle}}{\epsilon_{\text{sig}}(1 - P_{\text{decay}})} \frac{\sqrt{\sigma_{m^2}}}{\text{BR}(K^+ \to \pi^0 e^+ N; |U_e|^2 = 1)} \times \sqrt{\frac{\mathrm{d}N(K^+ \to \pi^0 e^+ \nu_e \gamma)}{\mathrm{d}m^2_{\text{miss}}}} \tag{3.3}$$

The median *projected exclusion* limit on $|U_e|^2$ from NA62 in the $K^+ \to \pi^0 e^+ N$ channel (valid for any number of quasi-degenerate HNLs) is presented in figure 3.4.1, along with the limits set by previous searches using the missing mass technique at KEK [168], PIENU [160] and NA62 [155], as well as the so-called seesaw "bounds" for both the normal and inverted hierarchy. These lines are the *model dependent* lower bounds on the mixing angle[3] $|U_e|^2 = \sum_{I=1,2} |\Theta_{eI}|^2$ in the type-I seesaw with two Majorana HNLs forming a quasi-Dirac pair (and assuming $P_{\text{decay}} \ll 1$). As discussed in section 3.3, we have assumed $\epsilon_{\text{bkg}} = 10^{-3}$ for in-acceptance photons, which results in an overall background efficiency of $\langle \epsilon_{\text{bkg}} \rangle \approx 1.7\%$ mainly driven by out-of-acceptance photons. If the HNL has visible decay channels and its lifetime is comparable to or smaller than the size of the detector, then the sensitivity to $|U_e|^2$ will be reduced by a factor of $(1 - P_{\text{decay}})^{-1}$ due to fewer events being available for the analysis. The approach considered here has no sensitivity to promptly decaying HNLs.

3.5 Discussion and outlook

The black solid line in figure 3.4.1 represents the sensitivity[4] achievable with the currently available dataset, corresponding to an effective number of K^+ decays of $N_K \approx 3 \times 10^{10}$. The NA62 collaboration is planning to collect an additional dataset in 2021 – 2024 [169]. Assuming no changes to the pre-scaling factors applied to the minimum-bias triggers, this leads to an estimated additional 6×10^{10} effective kaon decays. Considering in addition $K^+ \to \pi^0 e^+ N$ decays followed either by the Dalitz decay $\pi^0 \to \gamma e^+ e^-$ (which has branching fraction 1.17% [26]) or by a $\pi^0 \to \gamma\gamma$ decay with one of the photons converting just upstream of the trigger hodoscope, both of which are recorded by the current di-electron trigger [170], we expect an additional sample corresponding to 5×10^{10} kaon decays, bringing the total to 1.4×10^{11} by 2024. The corresponding sensitivity is shown by the black dashed line.

In order for the $K^+ \to \pi^0 e^+ N$ search at NA62 to become truly competitive in the region of interest, and start filling the current gap between 125 and 144 MeV, a dedicated trigger line (without the current prescaling factor of ~ 150) is required. If NA62 were to implement such a trigger for its 2021 – 2024 run, it would be able to establish a limit at the level of $|U_e|^2 \approx 10^{-7}$ (represented by the black dotted line) assuming a fully efficient trigger. Finally, any improvement in the rejection of out-of-acceptance photons, for instance through optimized selection or increased veto coverage, would push the sensitivity further down until the missed in-acceptance photons become the leading source of background.

The limits discussed in this paper present little model dependence, so long as the HNL is produced in a flavor-changing kaon decay. The remaining dependence comes from the possibly short lifetime of the HNL, which could induce additional activity in the detector when decaying, resulting in the event being excluded from the present analysis. In order to overcome this limitation, it would be interesting to allow for a displaced vertex compatible with the missing

[3] For consistency, we have plotted the lower bound on the mixing angle $|U_e|^2$ instead of the commonly used total mixing U^2. Our limit is therefore below the usual seesaw bound.

[4] The $(m_N,\ (1 - P_{\text{decay}})|U_e|^2)$ coordinates of the estimated sensitivity curves can be extracted from the file `figures/sensitivity.tex` in the LaTeX source of the arXiv version.

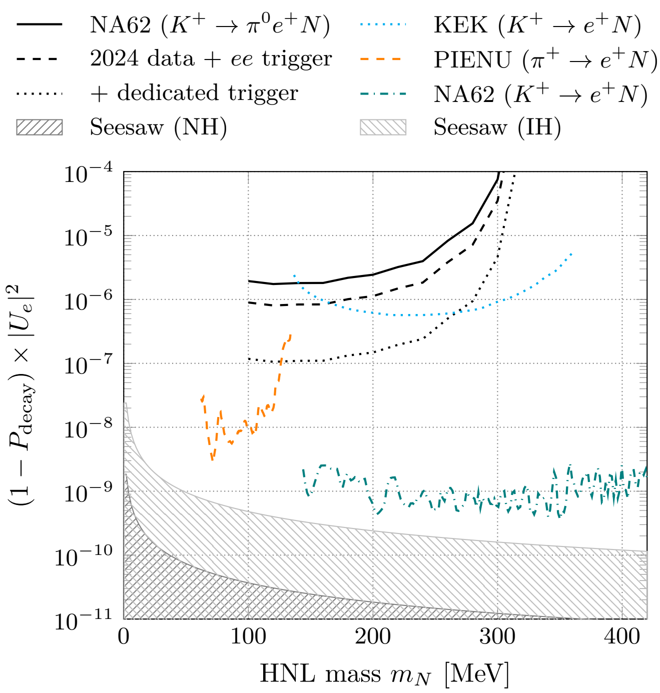

Figure 3.4.1: Projected exclusion reach of NA62 to HNLs in the $K^+ \to \pi^0 e^+ N$ channel (solid line), compared to the exclusion limits set by previous missing mass searches. The extra $1-P_{\text{decay}}$ factor corrects for the possibility that HNLs *visibly* decay inside the detector due to other types of interactions. Such events are not included in the present analysis, and would cause a weakening of the bound. For HNLs whose lifetime (in the laboratory frame) is significantly larger than the detector size, $P_{\text{decay}} \ll 1$ and the quantity probed is the usual mixing angle $|U_e|^2$. The seesaw "bounds" on $|U_e|^2$ are plotted under the assumption that two quasi-degenerate HNLs are fully responsible for neutrino oscillations (see the main text for details).

momentum. Finally, in order to probe the shorter lifetimes allowed by some non-minimal models (such as the one discussed in ref. [171]), dedicated searches involving prompt HNL decays will be needed. These searches are, however, inherently model dependent, since they target specific decay channels.

Acknowledgments

The authors are grateful to the NA62 Collaboration for allowing the use of their software code and framework to obtain part of the results presented here. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (GA 336581, 694896) and from the Carlsberg foundation. The work of IT has been supported by ERC-AdG-2015 grant 694896 and by the Swiss National Science Foundation Excellence grant 200020B_182864. We thank the anonymous referees for their thoughtful suggestions.

Chapter 4

Reinterpreting the ATLAS bounds on heavy neutral leptons in a realistic neutrino oscillation model

Authors: Jean-Loup Tastet, Oleg Ruchayskiy, Inar Timiryasov

In preparation

Abstract

Heavy neutral leptons (HNLs) are hypothetical particles, motivated in the first place by their ability to explain neutrino oscillations. Experimental searches for HNLs are typically conducted under the assumption of a single HNL mixing with a single neutrino flavor. Such exclusion limits do not directly constrain the corresponding mixing angles in *realistic HNL models* — those which can explain neutrino oscillations. The reinterpretation of the results of these experimental searches turns out to be a non-trivial task, that requires significant knowledge of the details of the experiment. In this work, we perform such a reinterpretation for the ATLAS search for promptly decaying HNLs in the tri-lepton final states. We show that in realistic HNL models, the actual limits may vary by several orders of magnitude, depending on the remaining free parameters of the model. Marginalizing over unknown model parameters leads to an exclusion limit on the total mixing angle which can be up to 3 orders of magnitude weaker than the limits reported in ref. [86]. This demonstrates that the reinterpretation of results from experimental searches is a *necessary* step to obtain meaningful limits on realistic models. We detail a few steps that can be taken by experimental collaborations in order to simplify the reuse of their results.

Contents

4.1 Introduction

4.1.1 Heavy neutral leptons

The idea that new particles need not be heavier than the electroweak scale, but rather can be light and feebly interacting is drawing increasing attention from both the theoretical and experimental communities, see e.g. [91, 106, 141, 142]. In particular, the hypothesis that heavy neutral leptons are responsible for (some of the) beyond-the-Standard-Model phenomena has been actively explored in recent years, see e.g. [35, 36, 49, 106, 141, 172, 173] and refs. therein. Heavy neutral leptons (HNLs) are massive particles that interact similarly to neutrinos, but with their interaction strength suppressed by flavor-dependent dimensionless numbers — *mixing angles* — $(U_e^2, U_\mu^2, U_\tau^2)$. HNLs first appeared in the context of left-right symmetric models [174–177] which required an extension of the fermion sector with Standard Model (SM) gauge singlet particles, and then in the (type I) see-saw mechanism [27–34] in which heavy Majorana neutrinos lead to light Standard Model neutrinos. The interest for these models increased when it was recognized that the same particles could also be responsible for the generation of the matter-antimatter asymmetry of the Universe [178]. This scenario (known as *leptogenesis*) has been actively developed since the 1980s (see reviews [179, 180]). In particular, it was found that the Majorana mass scale of right-handed neutrinos could be as low as the TeV, GeV or even MeV scale [35, 38, 45, 181–183]; for a recent overview see e.g. [184]. While two HNLs are sufficient to explain neutrino masses and oscillations as well as the origin of the matter-antimatter asymmetry, a third particle can play the role of dark matter [35, 36, 49, 50, 185] within the *Neutrino Minimal Standard Model* (νMSM).

Many experiments have searched for HNLs in the past (as summarized e.g. in refs. [91, 106, 163, 173, 186, 187]). Current generation particle physics experiments, including LHCb, CMS, ATLAS, T2K, Belle and NA62, all include HNL searches into their scientific programs [86, 97, 123–125, 155, 158, 188–193]. However, most of the existing or proposed analyses concentrate on the case of a *single HNL mixing with only one flavor*. Such a model serves as a convenient benchmark, but cannot explain any of the BSM phenomena that served as initial motivations for postulating HNLs. The same benchmarks are used when estimating the sensitivity of future experiments see e.g. [91], with an exception of the SHiP experiment, that provided sensitivity estimates for arbitrary sets of mixing angles [39]. This raises a few questions:

1. What HNL models explaining neutrino oscillations and/or other BSM phenomena are allowed or ruled out by previous searches? What parts of the HNL parameter space will be probed by future experiments?
2. What information do experimental groups need to provide in order to facilitate the answer to such questions in the future?

A number of tools exists, see e.g. [194–201] that allow for recasting of the LHC results for new sets of models, see also [202]. These tools have mostly been developed in the context of supersymmetry and similar searches at the LHC and are not readily applicable to the HNL models whose collider phenomenology is quite different.

In this work we perform a step in the direction of *recasting LHC results*. Specifically, we recast the ATLAS tri-lepton search [86] in the case of the *simplest realistic HNL model of neutrino oscillations*. This model features *two heavy neutral leptons* with (almost) degenerate masses. The possible values of the HNL mixings are constrained by neutrino oscillation data. In what follows we will refer to this model as a *realistic HNL model*.[1] As we shall see below, even in this simple model, the interpretation of the results is a non-trivial task.

4.1.2 Motivation for a reinterpretation

The realistic seesaw model describing neutrino oscillations brings several changes as compared to the single-HNL, single-flavor model analyzed by the ATLAS collaboration [86]. The analysis from ref. [86] concentrated on the following process:

$$pp \to W^\pm + X \quad \text{with} \quad W^\pm \to \ell_\alpha^\pm + N \quad \text{followed by} \quad N \to \ell_\alpha^\pm + \ell_\beta^\mp + \overset{(-)}{\nu}_\beta \tag{4.1}$$

where $\ell_\alpha^\pm$ are light leptons ($e^\pm$ or $\mu^\pm$), $\alpha \neq \beta$ and $\overset{(-)}{\nu}_\beta$ is a neutrino or anti-neutrino with flavor β. They have performed two independent analyses: one for the $e^\pm e^\pm \mu^\mp$+MET final state ("*electron channel*") and one for the $\mu^\pm \mu^\pm e^\mp$+ MET final state ("*muon channel*"). In both cases, only a single process (corresponding to diagram (b) in figure 4.1.1), along with its CP-conjugate, contributed to the final signal. The upper limit on an admissible signal was thus

[1] Notice that in the model with three HNLs the constraints on the mixings are much more relaxed [203, 204].

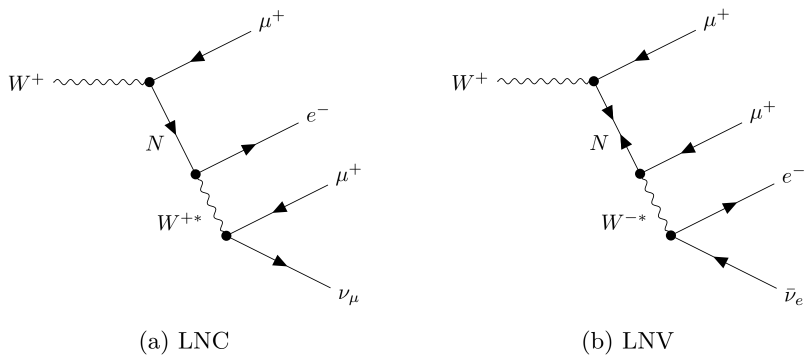

Figure 4.1.1: Lepton number conserving (LNC) and violating (LNV) diagrams contributing to the same $\mu^+\mu^+e^-$ + missing transverse energy (MET) final state.

directly translated into an upper bound on the mixing angle U_e^2 or U_μ^2, depending on the channel. The situation changes once we consider a realistic seesaw model with 2 HNLs:

1. In the two-HNLs model several processes contribute incoherently[2] to each final state. The upper bound on an admissible signal in any channel thus translates non-trivially into limits on all three mixings angles $(U_e^2, U_\mu^2, U_\tau^2)$.
2. Any set of mixing angles consistent with neutrino oscillation data leads to observable signals in both the $e^\pm e^\pm \mu^\mp$ and $\mu^\pm\mu^\pm e^\mp$ channels, therefore the statistical procedure should be changed and the predicted signal should be simultaneously fitted to both channels.
3. Different processes that contribute to the same tri-lepton final state have different kinematics (due in part to spin correlations [56]). Therefore the signal efficiencies need to be evaluated separately for every process.
4. We consider 2 HNLs with *nearly degenerate* masses. Due to HNL oscillations (cf. [56] or [63–66, 68, 69, 80] for earlier works) tiny mass differences (well below the mass resolution limit of ATLAS) can significantly affect the interference pattern, leading to the suppression or enhancement of some processes as compared to the single HNL case, see e.g. [48, 65, 111, 117]. For example, the kinematics of the two processes shown in figure 4.1.1 (and therefore their efficiencies) are different in the case of 1 and 2 HNLs *given the same set of mixing angles*. Thus the overall signal efficiency depends not only on the mixing angles, but also on the level of the HNL mass degeneracy.
   In order to account for this, we present our analysis for two limiting cases: the "Majorana-like" and "Dirac-like" (which we will define in section 4.2).

All these points make it impossible to reinterpret the ATLAS results by just rescaling them (as done e.g. in [85]). Instead one should perform a full signal and background modeling and evaluate the signal selection efficiencies. Although this can be properly done only by the collaboration itself, thanks to

[2] Their diagrams all produce different final states (when taking the light neutrino and its helicity into account) and therefore they do not interfere.

their access to the full detector simulation, the analysis framework and the actual counts in the signal regions, we will demonstrate that one can nonetheless reproduce the original ATLAS limits sufficiently well for the purpose of reinterpretation. Finally, we will briefly discuss what data from the collaboration could simplify our analysis and make it more precise, in the spirit of the LHC reinterpretation forum [202].

The present paper is organized as follows: In section 4.2 we introduce the notion of "realistic" seesaw models. To this end, we review the so-called type-I seesaw mechanism, discuss how neutrino oscillation data constrains its parameters, and examine how interference effects between multiple HNLs can completely change their phenomenology. We then describe our analysis procedure in section 4.3: we present the event selection, detail the calculation of the expected signal and efficiencies, and discuss our background model as well as the statistical method used to derive the exclusion limits. In section 4.4, we finally present our reinterpretation of the ATLAS limits on promptly-decaying HNLs within a realistic seesaw model with 2 HNLs, and we comment on these results. We conclude in section 4.5, and summarize what data should ideally be reported by experiments in order to allow reinterpreting their limits easily and accurately within different models.

4.2 Realistic neutrino oscillation models

4.2.1 The Lagrangian of the model

Our starting point is the type I seesaw mechanism [27, 28, 30, 31, 33, 34], that we briefly review below. The exposition is fairly standard and can be found, e.g. in refs. [49, 50, 106] and [26], ch. 14. The reader can skip it, taking notice of the definitions (4.3)–(4.4).

The Lagrangian of the model reads

$$\mathcal{L}_{\mathrm{SM+HNL}} = \mathcal{L}_{\mathrm{SM}} + i\bar{\nu}_{R_I}\partial\!\!\!/\nu_{R_I} - F_{\alpha I}(\bar{L}_\alpha \cdot \tilde{\Phi})\nu_{R_I} - \frac{1}{2}M_I\bar{\nu}^c_{R_I}\nu_{R_I}, \tag{4.2}$$

where $\mathcal{L}_{\mathrm{SM}}$ is the usual SM Lagrangian, ν_{R_I} are new, right-handed, particles that are SM gauge singlets. In the present paper we will consider the case of two HNLs, therefore the index I runs over 1, 2. L_α are the left-handed lepton doublets labeled with the flavor index $\alpha = e, \mu, \tau$ and $\tilde{\Phi} = i\sigma_2\Phi$, where Φ is the Higgs doublet. $F_{\alpha I}$ is the matrix of Yukawa couplings in the basis where the Yukawa couplings of charged leptons and the Majorana mass M_I of the right-handed neutrinos are both diagonal. After electroweak symmetry breaking, the Higgs field in the Lagrangian (4.2) obtains a vacuum expectation value $\langle\Phi\rangle = (0\, v)^T$ and the Yukawa interaction terms in eq. (4.2) effectively become Dirac mass terms coupling the left and right chiral components of the neutrinos. Since the right-handed neutrinos have, in addition, a Majorana mass, the spectrum of the theory is obtained by diagonalizing the full mass matrix.

For $|F_{\alpha I}v| \ll |M_I|$ one finds after the diagonalization 3 light mass eigenstates ν_i with masses m_1, m_2, m_3 and two heavy mass eigenstates N_I — the

HNLs — with masses M_1 and M_2.[3] As a consequence, the flavor eigenstates (SM neutrinos) $\nu_{L\alpha}$ can be expressed as a linear combination of the 5 mass eigenstates as

$$\nu_{L\alpha} = V^{\text{PMNS}}_{\alpha i}\nu_i + \Theta_{\alpha I} N_I^c\,, \tag{4.3}$$

where V^{PMNS} is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix (see e.g. [140]). As a result, the heavy mass eigenstates N_I contain an admixture of SM neutrinos $\nu_{L\alpha}$, and therefore possess "weak-like" interactions, suppressed by the mixing angles $\Theta_{\alpha I}$, approximately given by

$$\Theta_{\alpha I} \simeq \frac{vF_{\alpha I}}{M_I}\,. \tag{4.4}$$

4.2.2 Parametrization of the Yukawas

The Lagrangian (4.2) contains 11 new parameters, as compared to the SM one [106]. These parameters are, however, constrained by neutrino oscillation data [206]. Five neutrino parameters have already been measured: two mass differences (Δm^2_{atm} and Δm^2_{sun}) and three mixing angles ($\theta_{12}, \theta_{23}, \theta_{13}$). The remaining unknown parameters are the mass of the lightest neutrino, two Majorana phases, and the CP-violating phase δ. Our *a priori* choice of two HNLs restricts the mass of the lightest neutrino to be zero and only allows a certain combination of the Majorana phases to be independent. As a result, we are left with only two unknown parameters in the active neutrino sector, in addition to the discrete choice of the mass ordering.[4]

The measured low-energy parameters mean that for any choice of heavy neutrino masses M_I, the Yukawa couplings $F_{\alpha I}$ are not completely free. To account for this, we can parametrize the neutrino Yukawa couplings using the Casas-Ibarra parametrization [51]:

$$F = \frac{i}{v} V^{\text{PMNS}} \sqrt{m_\nu^{\text{diag}}}\, R \sqrt{M^{\text{diag}}}\,, \tag{4.5}$$

where the matrix $M^{\text{diag}} = \text{diag}\,(M_1, M_2)$, and R is a complex 3×2 matrix satisfying $R^{\text{T}} R = 1_{2\times 2}$. For the PMNS matrix we use the standard parametrization [26]. We parametrize the relevant combination of the Majorana phases in the PMNS matrix as $\eta = \frac{1}{2}(\alpha_{21} - \alpha_{31})$ for the normal neutrino mass hierarchy (NH), and $\eta = \frac{1}{2}\alpha_{21}$ for the inverted hierarchy (IH), with $\eta \in [0, 2\pi[$. The light neutrino mass matrix is $m_\nu^{\text{diag}} = \text{diag}(m_1, m_2, m_3)$ with $m_1 = 0$ for NH, and $m_2 = 0$ for IH.

In the model with two right-handed neutrinos, the matrices R depend on the neutrino mass hierarchy and are given by

$$R^{\text{NH}} = \begin{pmatrix} 0 & 0 \\ \cos\omega & \sin\omega \\ -\xi\sin\omega & \xi\cos\omega \end{pmatrix}, \qquad R^{\text{IH}} = \begin{pmatrix} \cos\omega & \sin\omega \\ -\xi\sin\omega & \xi\cos\omega \\ 0 & 0 \end{pmatrix} \tag{4.6}$$

[3] Given the Lagrangian (4.2) with two right-handed neutrinos, the lightest neutrino is massless (up to quantum corrections [205].

[4] These parameters may be probed in the not so distant future: for the inverted hierarchy, the next generation of neutrinoless double beta decay experiments may provide information on the Majorana phases [207], while the CP-violating phase δ is already constrained by T2K [208], with further improvements expected from the DUNE experiment [209].

with a complex angle $\omega = \mathrm{Re}\,\omega + i\,\mathrm{Im}\,\omega$, and a discrete parameter $\xi = \pm 1$. Changing the sign of ξ can be undone by $\omega \to -\omega$ along with $N_2 \to -N_1$ [53], so we fix $\xi = +1$.

4.2.3 Heavy neutrino mixing

The weak-like interactions of HNLs are suppressed by the *mixing angles* $\Theta_{\alpha I}$ defined in eq. (4.4). These mixing angles may contain complex phases, which play no role for the processes that we consider.[5] Only the cumulative effects of both N_1 and N_2 contributes to the observed signal and therefore the experimentally measurable quantities are

$$U_\alpha^2 \equiv \sum_I |\Theta_{\alpha I}|^2 \quad \text{and} \quad U_{\mathrm{tot}}^2 \equiv \sum_{\alpha, I} |\Theta_{\alpha I}|^2\,, \tag{4.7}$$

which respectively quantify the total HNL mixing to a particular flavor and the overall mixing between HNLs and neutrinos of definite flavor. The latter quantity has a particularly simple form in terms of the neutrino masses and Casas-Ibarra parameters:

$$U_{\mathrm{tot}}^2 = \frac{\sum_i m_i}{M_N} \cosh\left(2\,\mathrm{Im}\,\omega\right) \tag{4.8}$$

where $M_N = \frac{1}{2}(M_1 + M_2)$. For the corresponding expressions of U_α^2, see e.g. ref. [37].

As we have already mentioned, not all values of the Yukawa couplings $F_{\alpha I}$ — and hence of U_α^2 — are compatible with neutrino oscillation data. Only certain regions are allowed in $(U_e^2, U_\mu^2, U_\tau^2)$ space. For $|\,\mathrm{Im}\,\omega| \gg 1$ and $|M_1 - M_2| \ll M_N$, the shape of these regions *does not depend* on M_N, $|M_1 - M_2|$, or U_{tot}^2. Taking into account that

$$U_e^2/U_{\mathrm{tot}}^2 + U_\mu^2/U_{\mathrm{tot}}^2 + U_\tau^2/U_{\mathrm{tot}}^2 = 1\,, \tag{4.9}$$

we can display the combinations of U_α^2 which are compatible with neutrino oscillation data using a ternary plot as in figure 4.2.1, cf. [42, 85, 147]. In our analysis, we used the most recent global fit to neutrino oscillation data, NuFIT 5.0 [54, 55]. The shape of the allowed regions depends on the values of the Dirac phase δ and of the active neutrino mixing angle θ_{23}. We have used the three-dimensional projections of $\Delta\chi^2$ provided by NuFIT 5.0 in order to determine the 1, 2 and 3σ contours presented in figure 4.2.1. In order to better visualize the correspondence between the exclusion limits and various points in the allowed regions, we have defined a number of benchmarks, which are represented in figure 4.2.1.

4.2.4 Quasi-Dirac HNLs, lepton number violating effects and relevant limits

As neutrino oscillations do not constrain the masses of HNLs, M_1 and M_2 can be arbitrary. In this work we choose to consider the case where $M_1 \approx M_2$, i.e.

$$\Delta M \equiv |M_1 - M_2| \ll M_N = \frac{M_1 + M_2}{2}. \tag{4.10}$$

[5] These complex phases can be important if *the period of HNL oscillations* is comparable with the size of the experiment, see e.g. [56] and references therein.

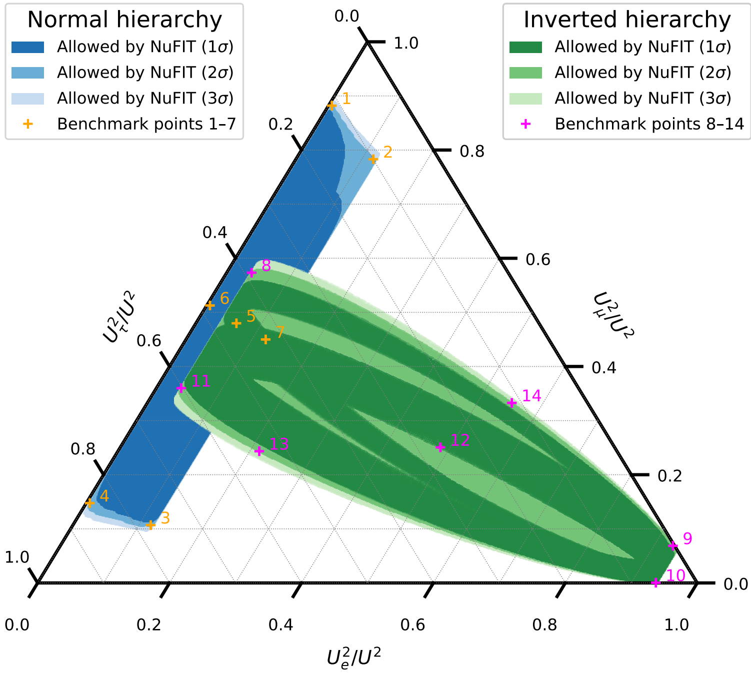

Figure 4.2.1: Ternary plot showing the combinations of mixing angles $U_\alpha^2/U_{\rm tot}^2$, $\alpha = e, \mu, \tau$, which are consistent with the NuFIT 5.0 [54, 55] fit to neutrino oscillation data, at the 1, 2 and 3σ levels, for the normal and inverted hierarchies. The markers denote the selected benchmark points, which are meant to represent both typical and extreme ratios of the squared mixing angles.

The motivation for this scenario is twofold. First, the mass degeneracy of two HNLs allows for sizable mixings between active neutrinos and HNLs in a technically natural way [48, 111, 114, 117, 210–215]. Secondly, low-scale leptogenesis (see the recent work [184] and references therein) requires a mass degeneracy between two heavy neutrinos. The mass splitting between the HNLs needs to be especially tiny if one wants to create the initial conditions required for the generation of sterile neutrino dark matter in the early Universe [38, 82, 185].

In the limit $M_1 \approx M_2$ there is an approximate global $U(1)$ symmetry in the theory.[6] In this *quasi-Dirac limit* of the two-HNLs model, the lepton number violating (**LNV**) processes (such as 4.1.1(b)) are suppressed compared to the lepton number conserving (**LNC**) processes. When $M_1 \neq M_2$ but $\Delta M \ll M_N$, HNL oscillations take place, as discussed in e.g. [35, 56, 63–66, 68, 69, 80]. As a result, lepton number violation may not be suppressed any more. Rather, the rates of LNC and LNV processes undergo a periodic modulation as a function of the proper time $\tau = \sqrt{(x_{\rm D} - x_{\rm P})^2}$ between the HNL production and decay

[6] The symmetry becomes exact when $M_1 = M_2$ and $\Theta_{\alpha 1} = \pm i \Theta_{\alpha 2}$. In this limit active neutrinos become massless and the two HNLs form a single Dirac particle Ψ such that $\frac{1+\gamma_5}{2}\Psi = \frac{\nu_{R_1} + i\nu_{R_2}}{\sqrt{2}}$.

vertices [56]:

$$d\Gamma^{\text{LNC/LNV}}_{\alpha\beta}(\tau) \cong 2\left|\Theta_{\alpha 1}\right|^2 \left|\Theta_{\beta 1}\right|^2 \Big(1 \pm \cos\left(\Delta M \tau\right)\Big) e^{-\Gamma\tau} d\widehat{\Gamma}^{\text{LNC/LNV}}_{\alpha\beta} \tag{4.11}$$

with the $(+)$ sign for LNC and $(-)$ for LNV, and where $d\widehat{\Gamma}^{\text{LNV/LNC}}_{\alpha\beta}$ is the differential rate for a tri-lepton process mediated by a single Majorana HNL N in the (unphysical) limit of a unit mixing angle between the HNL and the active flavor α at its production vertex, with flavor β at its decay vertex, and without the absorptive part; where $\Gamma \overset{\text{def}}{=} \Gamma_1 \cong \Gamma_2$ and by assumption $\Theta_{\alpha 2} \cong \pm i\Theta_{\alpha 1}$. Notice how in this quasi-Dirac limit, the oscillation pattern does not explicitly depend on the lepton flavors α and β, but only on whether the process is LNC or LNV. If ΔM vanishes exactly, then HNLs form a Dirac fermion and LNV effects are completely absent. Equation (4.11) demonstrates the two limiting cases of the two-HNLs seesaw model:

$$\begin{aligned} &\Delta M\tau \ll 2\pi \quad \text{(Dirac-like limit)} && d\Gamma^{\text{LNV}}_{\alpha\beta} \approx 0,\ d\Gamma^{\text{LNC}}_{\alpha\beta} \text{ is enhanced by } \cong 2 \\ &\Delta M\tau \gg 2\pi \quad \text{(Majorana-like limit)} && \text{integrated partial widths } \Gamma^{\text{LNV}}_{\alpha\beta} \cong \Gamma^{\text{LNC}}_{\alpha\beta} \end{aligned} \tag{4.12}$$

where τ must satisfy both $\tau\Gamma_N \lesssim 1$ and $\gamma\tau \lesssim L_{\text{det}}$ (whichever is stronger), with Γ_N denoting the total HNL width, γ its boost factor, and L_{det} the typical detector size.

In this work we will consider these two limiting cases for quasi-Dirac HNLs:

- **Dirac-like**: the pure Dirac ($\Delta M = 0$) limit where all LNV effects are completely absent, and LNC rates are coherently enhanced by a factor of 2;
- **Majorana-like**: the $\Delta M\tau \gg 2\pi$ limit where both LNV and LNC processes are present, with the same integrated rates.

Comparing these two limiting cases for the same benchmark models allows to assess the level of uncertainty introduced by the unknown ΔM.

4.3 Procedure

In order to reinterpret the limits from the ATLAS prompt search [86] (with extra details in the Ph.D. thesis [87]) we have tried to reproduce the ATLAS analysis as accurately as possible. Our signal is simulated using MadGraph5_aMC@NLO [216] with the HeavyN model [217,218] (section 4.3.2). For the event selection (section 4.3.1), we have implemented the ATLAS cut flow and obtained comparable efficiencies (section 4.3.3). We take the total background counts from the ATLAS publication [86] (section 4.3.4). Finally, in order to compute the limits (section 4.3.5), we use the CL_s test statistics, along with a very simplified treatment of uncertainties.

4.3.1 Event selection

The prompt ATLAS analysis [86] considers the final states consisting of three isolated charged leptons (with electron or muon flavor) with no opposite-charge same-flavor lepton pairs (in order to limit the background from Z decays), i.e. only $e^{\pm}e^{\pm}\mu^{\mp}$ (electron channel) and $\mu^{\pm}\mu^{\pm}e^{\mp}$ (muon channel) are considered. It focuses on HNLs which are sufficiently short-lived that their decay vertex can be efficiently reconstructed using the standard ATLAS tracking algorithm. Since our reinterpretation will include a number of processes not included in the original ATLAS analysis and having different kinematics (e.g. LNC processes, which are absent in the single-flavor mixing assumption), we cannot use the published ATLAS efficiencies and we have to compute them on our own.

As we will see, imposing the same *analysis cuts* only is sufficient to accurately reproduce the ATLAS efficiencies. This seems to indicate that these cuts are stronger than the requirements imposed by trigger, tracking, etc. Therefore we need not worry about the technical details of the experiment or detailed *detector simulations*, and we can just focus on the "*cut flow*". The list of cuts is shown in table 4.3.1, and their order roughly follows that of ref. [87].

1. We start by applying a cut on the distance of closest approach to the origin in the r–z plane, i.e. $|\Delta z_0 \sin(\theta)| < 1\,\mathrm{mm}$ for all three leptons.
2. We then apply the transverse momentum and pseudorapidity requirements on the three changed leptons, i.e. $p_{\mathrm{T}} \geq 4.5\,\mathrm{GeV}$ and $|\eta| \in [0, 1.37[\cup]1.52, 2.47]$ for all electrons[7] and $p_{\mathrm{T}} \geq 4\,\mathrm{GeV}$ and $|\eta| \leq 2.5$ for all muons.
3. Next, we apply a p_{T}-dependent weight to each electron in order to simulate the efficiency of the lepton identification requirement (using the "loose" working point). We use for that the p_{T}-differential efficiency reported in ref. [219]. For muons, the efficiency is close to 1 [220], so we do not apply any weight.
4. Next, we require the tri-lepton invariant mass M_{3l} to be in the interval $]40, 90[\,\mathrm{GeV}$.
5. Next, we apply the trigger offline requirements on the two leading leptons, i.e. $p_{\mathrm{T}}(e_{\mathrm{lead}}) > 27\,\mathrm{GeV}$ and $p_{\mathrm{T}}(e_{\mathrm{sublead}}) > 10\,\mathrm{GeV}$ for the electron channel and $p_{\mathrm{T}}(\mu_{\mathrm{lead}}) > 23\,\mathrm{GeV}$ and $p_{\mathrm{T}}(\mu_{\mathrm{sublead}}) > 14\,\mathrm{GeV}$ for the muon channel.
6. Next we apply a weight to each lepton in order to simulate the efficiency of lepton isolation. We use the p_{T}-differential isolation efficiencies reported in ref. [219] for electrons and ref. [220] for muons, using the "loose" working point in both cases.
7. For the electron channel only, a further cut is applied on the invariant mass of the $e^{\pm}e^{\pm}$ pair, $M(e,e) < 78\,\mathrm{GeV}$, in order to veto the background from $Z \to e^+e^-$ where one of the electron charges is misreconstructed.
8. Finally, the missing transverse energy is restricted to $E_{\mathrm{T}}^{\mathrm{miss}} < 60\,\mathrm{GeV}$.

Our cut flow is summarized in table 4.3.1. One notable difference from the ATLAS paper is the absence of a b-jet veto in our analysis, which we omitted since it almost does not affect the signal. A further difference comes from the cuts related to the displacement of the leading lepton. ATLAS imposes $|\Delta z_0 \sin(\theta)|(l_{\mathrm{lead}}) < 0.5\,\mathrm{mm}$ and $|d_0/\sigma(d_0)| < 5$ (electron) or < 3 (muon), while we impose $|\Delta z_0 \sin(\theta)|(l) < 1\,\mathrm{mm}$ on all leptons[8] and omit the d_0 cut. This

[7] We use the 2016 values here.

[8] We followed ref. [87] here, but we plan to change our cuts to match ref. [86] in the final publication.

#	Electron channel	Muon channel
1	$\lvert\Delta z_0 \sin(\theta)\rvert(l) < 1\,\mathrm{mm}$	
2	$p_\mathrm{T}(e) \geq 4.5\,\mathrm{GeV}$, $p_\mathrm{T}(\mu) \geq 4\,\mathrm{GeV}$ $\lvert\eta(e)\rvert \in [0, 1.37[\cup]1.52, 2.47]$, $\lvert\eta(\mu)\rvert \leq 2.5$	
3	"Loose" electron ID	
4	$40\,\mathrm{GeV} < M(l, l, l') < 90\,\mathrm{GeV}$	
5	$p_\mathrm{T}(e_\mathrm{lead}) > 27\,\mathrm{GeV}$ $p_\mathrm{T}(e_\mathrm{sublead}) > 10\,\mathrm{GeV}$	$p_\mathrm{T}(\mu_\mathrm{lead}) > 23\,\mathrm{GeV}$ $p_\mathrm{T}(\mu_\mathrm{sublead}) > 14\,\mathrm{GeV}$
6	"Loose" lepton isolation	
7	—	Z veto: $M(e, e) < 78\,\mathrm{GeV}$
8	$E_\mathrm{T}^\mathrm{miss} < 60\,\mathrm{GeV}$	

Table 4.3.1: Our cut flow for the electron and muon channels.

Electron channel ($e^\pm e^\pm \mu^\mp$)				
Process	ΔL	α	β	MadGraph process string
$W^+ \to e^+(N \to \mu^- e^+ \nu_e)$	0	e	μ	`p p > e+ n1, n1 > mu- e+ ve`
$W^- \to e^-(N \to \mu^+ e^- \bar{\nu}_e)$	0	e	μ	`p p > e- n1, n1 > mu+ e- ve~`
$W^+ \to e^+(N \to e^+ \mu^- \bar{\nu}_\mu)$	-2	e	e	`p p > e+ n1, n1 > e+ mu- vm~`
$W^- \to e^-(N \to e^- \mu^+ \nu_\mu)$	$+2$	e	e	`p p > e- n1, n1 > e- mu+ vm`

Table 4.3.2: Signal processes contributing to the electron channel.

does not affect the signal since the leading lepton has a very small displacement $\ll 1\,\mathrm{mm}$ in all relevant cases.[9]

Finally, the events passing the above cuts are binned in $M(l_\mathrm{sublead}, l')$, which approximates the invariant mass of the HNL for small HNL masses (for which the leading lepton is usually the prompt lepton). The bins are $[0, 10[$, $[10, 20[$, $[20, 30[$, $[30, 40[$ and $[40, 50[\,\mathrm{GeV}$.

4.3.2 Signal

In order to reinterpret the sensitivity of the ATLAS prompt HNL search for arbitrary combinations of HNL masses M_N and ratios of mixing angles, we need to be able to compute the expected signal counts in each $M(l_\mathrm{sublead}, l')$ bin in each signal region, for any model parameters. We do so using a simple model, described below.

MadGraph setup

The signal processes contributing to each channel are listed in tables 4.3.2 and 4.3.3.[10] For Majorana-like HNL pairs, all processes contribute, while for Dirac-like HNL pairs only those which conserve the total lepton number ($\Delta L = 0$) contribute (with a factor-of-2 enhancement for the total cross section).

[9] For light HNLs, the leading lepton is almost always the prompt lepton from the W decay, while heavier HNLs decay with a very short displacement due to their much shorter lifetime.

[10] In the "Process" column, we use a bar to indicate the chirality of the produced light neutrinos. Their Majorana nature does not play a role here.

Muon channel ($\mu^\pm\mu^\pm e^\mp$)				
Process	ΔL	α	β	MadGraph process string
$W^+ \to \mu^+(N \to e^-\mu^+\nu_\mu)$	0	μ	e	`p p > mu+ n1, n1 > e- mu+ vm`
$W^- \to \mu^-(N \to e^+\mu^-\bar{\nu}_\mu)$	0	μ	e	`p p > mu- n1, n1 > e+ mu- vm~`
$W^+ \to \mu^+(N \to \mu^+e^-\bar{\nu}_e)$	-2	μ	μ	`p p > mu+ n1, n1 > mu+ e- ve~`
$W^- \to \mu^-(N \to \mu^-e^+\nu_e)$	$+2$	μ	μ	`p p > mu- n1, n1 > mu- e+ ve`

Table 4.3.3: Signal processes contributing to the muon channel.

For each process, we generate a Monte-Carlo sample which will be used to compute both the cross section and the efficiency. Each sample consists of $\sim$ 40000 weighted events generated at leading order using MadGraph5_aMC@NLO v2.8.x [216] along with the HeavyN model [217, 218] (specifically, we use the SM_HeavyN_NLO model). The center of mass energy is set to $\sqrt{s} = 13\,\text{TeV}$ and the integrated luminosity to $\mathcal{L}_{\text{int}} = 36.1\,\text{fb}^{-1}$, in order to match the parameters of the 2019 prompt analysis. We generate the processes listed in the "MadGraph process string" column in tables 4.3.2 and 4.3.3, with up to two additional hard jets. Pythia 8 is then used (through the MadGraph interface) to shower and hadronize the events. We use the event weights and the merged cross section reported by Pythia.

Signal computation for arbitrary model parameters

In order to obtain the physical cross section, a number of model parameters need to be specified: the HNL mass M_N, its mixing angles[11] $|\Theta_e|$, $|\Theta_\mu|$ and $|\Theta_\tau|$ and its total decay width Γ_N. Generating a new sample for every set of parameters would be computationally prohibitive. Fortunately, we can leverage the scaling properties of the cross section in order to exactly recompute the cross section to each new sets of mixing angles. This is done as follows.

As a first step, we generate Monte-Carlo samples for all the processes listed in tables 4.3.2 and 4.3.3, for each HNL mass $M_N \in \{5, 10, 20, 30, 50\}\,\text{GeV}$ and using the reference parameters $|\Theta|_{\text{ref}} = 10^{-3}$ and $\Gamma_{\text{ref}} = 10^{-5}\,\text{GeV}$ as placeholders for the remaining model parameters.[12] For each process P, we only set the relevant mixing angle $|\Theta_{\alpha(P)}|$ and $|\Theta_{\beta(P)}|$ to $|\Theta|_{\text{ref}}$, where $\alpha(P)$ and $\beta(P)$ respectively correspond to the generations coupling to the HNL at production and decay, as listed in tables 4.3.2 and 4.3.3.

The key observation here is that the branching fraction of $W^+ \to l_\alpha N$ is proportional to $|\Theta_\alpha|^2$, while the branching fraction of $N \to l_\beta l_\gamma \nu_\gamma$ is proportional to $|\Theta_\beta|^2/\Gamma_N$. Therefore, the cross section for a given process P is proportional to $|\Theta_{\alpha(P)}|^2|\Theta_{\beta(P)}|^2/\Gamma_N$. Starting from the reference cross section σ_P^{ref} obtained for the reference parameters, this allows to extrapolate the physical cross section to new parameters:

$$\sigma_P(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) = \sigma_P^{\text{ref}} \times \frac{|\Theta_{\alpha(P)}|^2|\Theta_{\beta(P)}|^2}{|\Theta|^4_{\text{ref}}} \times \frac{\Gamma_{\text{ref}}}{\Gamma_N(M_N, \Theta_e, \Theta_\mu, \Theta_\tau)} \tag{4.13}$$

[11] Since we are dealing with 2 HNLs far from the seesaw line, $\Theta_{\alpha 2} \cong \pm i\Theta_{\alpha 1}$ [48, 111]. We generate the Monte-Carlo samples for a single HNL with parameters $\Theta_\alpha \stackrel{\text{def}}{=} \Theta_{\alpha 1}$, such that $|\Theta_\alpha| = |\Theta_{\alpha 1}| \cong |\Theta_{\alpha 2}| = \frac{1}{2}U_\alpha^2$, see eq. (4.7).

[12] These parameters allow for the successful numerical integration in the narrow width approximation.

Nature	$c_P,\ P \in \mathrm{LNC}$	$c_P,\ P \in \mathrm{LNV}$	$c_\Gamma = \Gamma_N/\Gamma_{\mathrm{Maj.}}$
One Majorana HNL (reference)	**1**	**1**	**1**
One Dirac HNL	1	0	1/2
Quasi-Dirac pair: **Majorana-like**	2	2	1
Quasi-Dirac pair: **Dirac-like**	4	0	1

Table 4.3.4: Multiplicative coefficients c_P to be applied to the cross section of each process P, and c_Γ to be applied to the total HNL width Γ_N, depending on the HNL(s) nature and on whether the process is LNC or LNV.

Since the total HNL width enters this formula, we need to be able to compute it for arbitrary parameters too. To this end we follow a similar approach. We notice that the partial width into a given decay channel D is proportional to $|\Theta_{\beta(D)}|^2$, where $\beta(D)$ denotes the flavor with which the HNL mixes when decaying. Summing over all decay channels and all three flavors, we can then express the total decay width as:

$$\tau_N^{-1} = \Gamma_N(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) = \sum_{\beta=e,\mu,\tau} |\Theta_\beta|^2 \times \hat{\Gamma}_\beta(M_N) \tag{4.14}$$

where $\hat{\Gamma}_\beta(M_N) = \Gamma_N(M_N, \delta_{\beta e}, \delta_{\beta\mu}, \delta_{\beta\tau})$ is the total decay width obtained by setting $\Theta_\beta = 1$ and the two other mixing angles to zero. It can be easily computed with MadGraph by generating the `n1 > all all all` process. This extrapolation method, which makes use of the scaling properties of the relevant branching fractions, has been successfully validated by explicitly computing the cross section for a few non-trivial benchmark points and comparing the results. The contribution N_P of a given process P to the total event count (before applying any selection) is then obtained by multiplying the relevant cross section by the integrated luminosity: $N_P = \mathcal{L}_{\mathrm{int}} \times \sigma_P$.

Signal computation for quasi-Dirac HNLs

Finally, since the signal samples have been computed for a single Majorana HNL, we need to apply a correction factor c_P to each cross section when considering a quasi-Dirac HNL pair. If this HNL pair is Majorana-like (i.e. it has both LNC and LNV processes with equal rates), then all cross sections must be multiplied by 2, since there are two mass eigenstates whose event rates add incoherently. However, for a Dirac-like HNL pair (which only has LNC processes), the LNC cross sections must be multiplied by 4 due to the coherent enhancement discussed in section 4.2.4, while the LNV ones should all be set to zero. Unlike in the case of a single Dirac fermion, no correction to the total HNL width needs to be applied. The correction factors are summarized in table 4.3.4.

4.3.3 Efficiencies

In order to obtain a sensitivity estimate, we must compute the expected signal count in every $M(l_{\mathrm{sublead}}, l)$ bin reported by the ATLAS collaboration.[13]. This

[13] We consider both signal regions (for the $e^\pm e^\pm \mu^\mp$ and $\mu^\pm \mu^\pm e^\mp$ signatures) simultaneously, so there are 10 bins in total: 5 in the electron channel and 5 in the muon channel.

is done by multiplying the true signal count by a signal efficiency. Since the relative contributions of the various diagrams — which all have different kinematics and therefore different efficiencies — depend on the model parameters, in general we expect the signal efficiency to depend on the mass M_N, nature (Majorana-like or Dirac-like), lifetime τ_N and all the mixing angles of the quasi-Dirac HNL pair. However, when considering a single process / diagram, the nature and mixing angles "factor out" such that the efficiency for this process depends only on the mass and lifetime. We therefore need to compute one efficiency $\epsilon_{P,b}(M_N, \tau_N)$ for every process P and every bin b. The total event count in bin b is then computed by summing over all the processes:

$$N_b = \mathcal{L}_{\text{int}} \times \sum_P \epsilon_{P,b}(M_N, \tau_N) \times c_P \times \sigma_P(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) \tag{4.15}$$

where c_P is the correction factor applied to the cross section for quasi-Dirac HNLs.

For a given process P and bin b, the efficiency $\epsilon_{P,b}(M_N, \tau_N)$ is computed by filtering the corresponding Monte-Carlo sample through the cut flow described in section 4.3.1 and table 4.3.1. The binned efficiency is then:

$$\epsilon_{P,b} = \frac{\sum(\text{weights of events after cuts, which end up in bin } b)}{\sum(\text{weights of all events before cuts, from any bin})} \tag{4.16}$$

where the sums run over all events generated for the process P and the events which fail to pass a given cut have their weight set to zero.[14] Similarly, we can obtain the unbinned efficiency as:

$$\epsilon_P = \frac{\sum(\text{all event weights after cuts})}{\sum(\text{all event weights before cuts})}. \tag{4.17}$$

The unbinned efficiencies for the four LNV processes are plotted in figure 4.3.1 along with the efficiencies reported for ATLAS in ref. [87], while those for LNC processes are plotted in figure 4.3.2. Since the efficiency of a process depends on both the HNL mass and its lifetime, we had to choose some benchmark points to produce figures 4.3.1 and 4.3.2. In order to be able to compare our efficiency calculation with the ATLAS efficiencies, we have chosen the same benchmarks as reported in ref. [87] and reproduced in table 4.3.5. Our estimate is accurate for three LNV processes (with a mean relative error of 0.07), while for $W^- \to e^-(N \to e^-\mu^+\nu_\mu)$ the relative error can be as high as 0.5. Error bars denote the rounding error of $\pm 5 \times 10^{-3}$ in the ATLAS efficiencies listed in ref. [87], and the missing entries in the electron channel correspond to efficiencies which have been rounded down to zero. Comparing figure 4.3.2 with figure 4.3.1, notice how **the efficiencies for LNC processes are measurably smaller than for LNV processes**, sometimes by up to a factor of ~ 2. This is mostly due to the different spin correlation patterns for LNC vs. LNV leading to different lepton p_T spectra and to different geometrical acceptances of the lepton p_T cuts.

Even using the extrapolation method described above and eq. (4.15), one efficiency $\epsilon_{P,b}(M_N, \tau_N)$ must in principle still be computed for every process P,

[14] Some cuts (such as lepton ID and isolation cuts) are implemented by reweighting events using tabulated efficiencies.

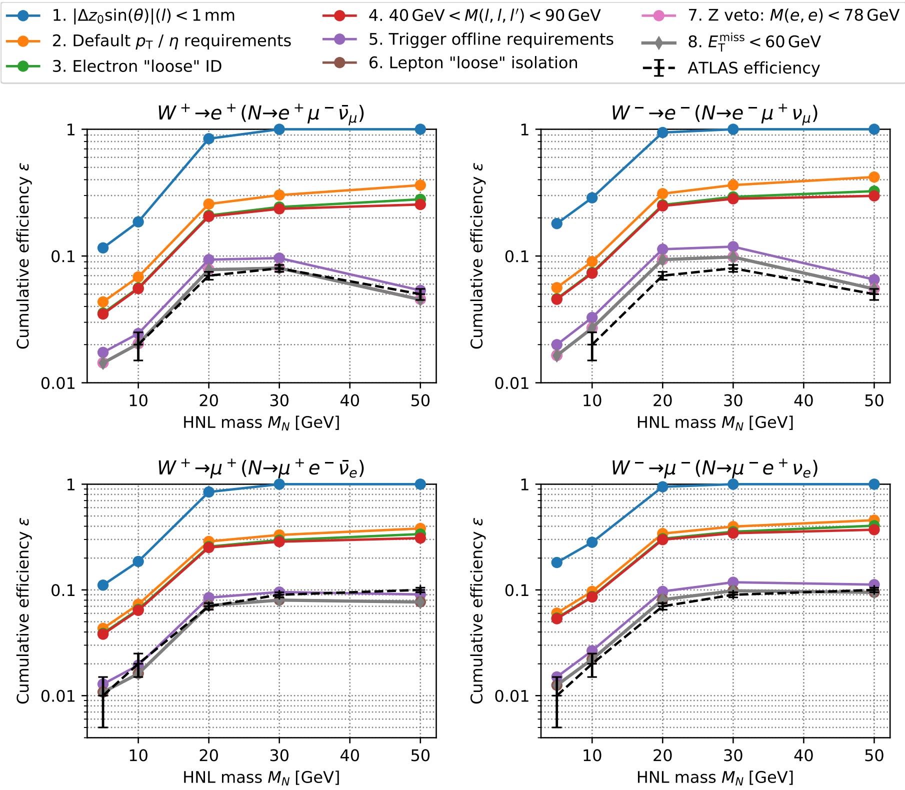

Figure 4.3.1: Cumulative unbinned signal efficiencies (for the total event count, i.e. summed over all bins) after applying each cut listed in table 4.3.1, computed for the benchmark points found in ref. [87]. The black dashed line denotes the total efficiencies reported in ref. [87], and should be compared to the gray line with diamond markers (which corresponds to all cuts being applied). These efficiencies are for lepton number *violating* (LNV) processes only, since these were the relevant processes in the original prompt search.

HNL mass M_N	5 GeV	10 GeV	20 GeV	30 GeV	50 GeV
HNL lifetime τ_N	1 mm	1 mm	0.1 mm	0.01 mm	1 µm

Table 4.3.5: Benchmark points (taken from ref. [87]) used to plot the efficiencies in figures 4.3.1 and 4.3.2. Note that our calculation is more general, and works for any combination of M_N and τ_N.

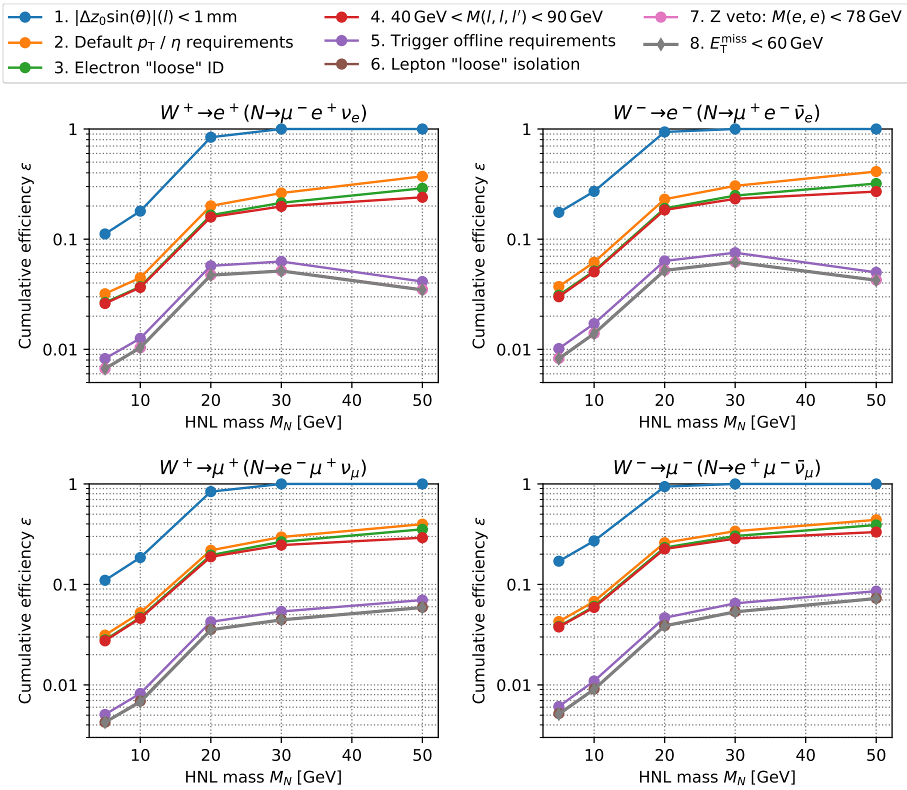

Figure 4.3.2: Cumulative unbinned signal efficiencies (for the total event count, i.e. summed over all bins) after applying each cut listed in table 4.3.1, computed for the benchmark points found in ref. [87], for lepton number *conserving* (LNC) processes. The gray line with diamond markers corresponds to the total efficiency.

bin b, HNL mass M_N and lifetime τ_N. However, several simplifications exist. First, the efficiencies for the full set of $M(l_{\mathrm{sublead}}, l')$ bins (keeping the other parameters fixed) can be computed simultaneously, since the events only need to go through the cut flow once, before the binning is applied. More interestingly, it also turns out that the τ_N dependence can be quite accurately parametrized using a simple functional form $\epsilon(\tau_N)$. This functional form can be constrained by requiring the following asymptotic behavior:

- $\epsilon(\tau_N) \to \epsilon_0$ (prompt efficiency) as $\tau_N \to 0$.
- $\epsilon(\tau_N) \propto \frac{1}{\tau_N}$ for sufficiently large τ_N.

The "simplest" functional form satisfying these two conditions is:

$$\epsilon(\tau_N) = \frac{\epsilon_0}{1 + \frac{\tau_N}{\tau_0}} \tag{4.18}$$

with ϵ_0 the prompt efficiency and τ_0 the typical lifetime after which the efficiency starts to drop due to the HNL displacement. After fitting it to the efficiencies which have been explicitly computed for a number of lifetime points,

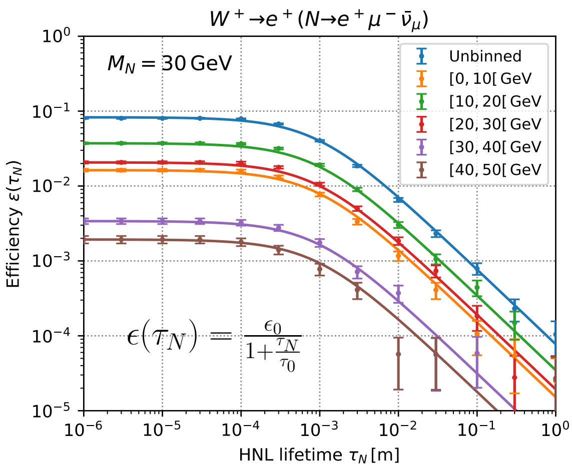

Figure 4.3.3: Binned and unbinned efficiencies as a function of the HNL lifetime τ_N, for the process $W^+ \to e^+(N \to e^+\mu^-\bar{\nu}_\mu)$ with $M_N = 30\,\mathrm{GeV}$. The dots represent the efficiencies calculated explicitly, while the lines correspond to the fitted model. Error bars denote an *estimate* of the statistical uncertainties from the finite size of the Monte-Carlo sample.

this model can be used to extrapolate the efficiency to arbitrary HNL lifetimes. As an example, the model, along with the lifetime points used for the fit, are presented in figure 4.3.3 for both the binned and unbinned efficiencies, for the $W^+ \to e^+(N \to e^+\mu^-\bar{\nu}_\mu)$ process with a 30 GeV HNL. The relative error between the data and the model is $\lesssim 10\%$ (on top of the statistical error). The efficiencies for other processes and mass points display a similar behavior.

Thanks to these simplifications, for each HNL mass M_N and process P, the efficiencies need only be computed for 3 or more lifetime points in order to obtain the full lifetime dependence along with an error estimate. This amounts to 12 or more Monte-Carlo samples per mass point for Dirac-like HNL pairs, and 24 or more for Majorana-like HNL pairs.[15] This makes the approach computationally tractable (although expensive) for experiments who would like to report their efficiencies in a benchmark-agnostic way, while still using their full detector simulation.

4.3.4 Background

A number of Standard Model processes can mimic the signatures that we are looking for. This can happen if these processes have the same final state (irreducible background) or if they are misreconstructed as the same final state (reducible background) due to *fake* leptons (i.e. non-prompt leptons from jets or leptons from pileup). ATLAS has found the irreducible background to be subdominant [86], and the main background components to be *multi-fakes* (multiple fake leptons coming from W+jets or multiple jets) as well as $t\bar{t}$ with a fake lepton.

[15] Plus three samples for computing the HNL lifetime, but these only need to be run at parton level and therefore have a negligible computational cost.

Each of these background sources comes with statistical uncertainties. The kinematic distribution of the multi-fake sample is estimated from data using a number of *estimation* regions, then normalized by fitting a normalization factor μ_{mf} to the three control regions. Due to the finite sizes of the data samples, both of these steps introduce statistical errors into the multi-fake estimate, with potentially non-trivial correlations between the $M(l_{\text{sublead}}, l')$ bin counts, which we are ultimately interested in. Similarly, the finite size of the $t\bar{t}$ Monte-Carlo sample and the finite event counts in the control regions used to estimate its normalization factor $\mu_{t\bar{t}}$ also introduce statistical errors into the $t\bar{t}$ estimate.

The detailed uncertainties (including correlations) of the individual background components are not listed in ref. [86]. Performing a detailed background analysis is out of the scope of the present paper. Instead, we have decided to use a simplified background model, which only takes into account the total background count in each bin, but is nonetheless capable of providing a good enough approximation of the sensitivity for the purpose of this reinterpretation.

To this end, the total background count in each channel and each $M(l_{\text{sublead}}, l')$ bin, along with its uncertainty band, is digitized from figure 5 in ref. [86]. Since the statistical errors on the normalization factors μ_{mf} and $\mu_{t\bar{t}}$ are among the leading uncertainties, we will assume that the uncertainty bands reported in this figure are entirely caused by a single normalization factor μ_{tot}, or in other words that the uncertainties in the various $M(l_{\text{sublead}}, l')$ bins are maximally correlated. The accuracy of this simplified model will be explicitly tested in section 4.3.5.

4.3.5 Statistical limits

Ref. [86] found a very good compatibility between the observed counts and the background-only hypothesis. They then proceeded with exclusion limits by testing the compatibility of the observed counts under the signal + background hypotheses for five different benchmark points in the (mass, lifetime) space, each for two different mixing patterns: with electron or muon flavor.

In order to define the exclusion limit, ATLAS uses the CL_s test [221]. For completeness, a quick reminder about the CL_s technique follows in section 4.3.5. Knowledgeable users are welcome to skip it and go directly to section 4.3.5.

CL_s technique: a general reminder

The CL_s technique is based on the likelihood-ratio test statistics, more specifically on:

$$t(x) \equiv 2 \ln \left(\frac{\mathcal{L}(x|H_{s+b})}{\mathcal{L}(x|H_b)} \right) \tag{4.19}$$

where $\mathcal{L}$ denotes the likelihood, x the data, H_b the background-only hypothesis and H_{s+b} a signal + background hypothesis. Larger values of t indicate more signal-like data. The distribution of t is estimated under each hypothesis through the use of pseudo-experiments X: $p_b(t) = \mathcal{P}(t(X))$ for $X \sim H_b$ and $p_{s+b}(t) = \mathcal{P}(t(X))$ for $X \sim H_{s+b}$. Given an observation x_{obs} and the corresponding value of the test statistics $t_{\text{obs}} = t(x_{\text{obs}})$, the CL_b and CL_{s+b} values are then computed as:

$$\mathrm{CL}_b = \mathcal{P}\left(t(X) < t_{\mathrm{obs}} | H_b\right) = \int_{-\infty}^{t_{\mathrm{obs}}} \mathrm{d}t\, p_b(t) \tag{4.20}$$

$$\mathrm{CL}_{s+b} = \mathcal{P}\left(t(X) < t_{\mathrm{obs}} | H_{s+b}\right) = \int_{-\infty}^{t_{\mathrm{obs}}} \mathrm{d}t\, p_{s+b}(t) \tag{4.21}$$

In other words, CL_b and CL_{s+b} are the probabilities of obtaining a dataset that is more background-like than the observed one, respectively under the background and signal + background hypotheses. Both increase for increasingly signal-like x_{obs}. Finally, the value of the CL_s test statistics is given by the ratio:

$$\mathrm{CL}_s = \frac{\mathrm{CL}_{s+b}}{\mathrm{CL}_b} \tag{4.22}$$

and a given signal + background hypothesis H_{s+b} is considered to be excluded if $\mathrm{CL}_s < 0.05$. For any signal stronger than the $\mathrm{CL}_s = 0.05$ limit, the probability of a type-I error (false exclusion) will always be less than 0.05. In order to complete the statistical analysis, the likelihood remains to be specified. We will proceed with this in the following section.

CL_s technique: implementation

The observables in question are the event counts in the two signal regions (for the electron and muon channels), each channel consisting of 5 $M(l_{\mathrm{sublead}}, l')$ bins. Since we will be dealing with non-trivial combinations of mixing angles, we simultaneously include both channels in our likelihood. We thus end up with 10 bin counts $\{x_i\}$, with $i = 1 \ldots 5$ for the electron channel and $i = 6 \ldots 10$ for the muon channel. As discussed in section 4.3.4, we model the background as a set of expectation values $\{b_i\}$ for each bin $i = 1 \ldots 10$ (taken from the ATLAS paper) along with a Gaussian-constrained normalization factor μ_{tot} with standard deviation $\sigma_{\mathrm{tot}} = \sum_i (b_i^+ - b_i^-)/(2\sum_i b_i)$, where the $^-$ and $^+$ superscripts respectively denote the lower and upper uncertainty bands from the ATLAS plot (see table 4.3.6). The signal is modeled as a set of signal expectations $\{s_i\}$, $i = 1 \ldots 10$, which we compute for each set of the model parameters $(M_N, \Theta_e, \Theta_\mu, \Theta_\tau)$ using the method described in sections 4.3.2 and 4.3.3. Contrary to ATLAS, we do not use a signal strength parameter μ, since this would amount to rescaling the mixing angles without changing the lifetime, leading to inconsistent results.[16] We neglect all uncertainties on the signal counts, which we have estimated to be at the sub-percent level. The bin counts x_i are assumed to be Poisson distributed, with expectation values of respectively $\mu_{\mathrm{tot}} b_i$ for the background-only hypothesis and $\mu_{\mathrm{tot}} b_i + s_i$ for the signal + background

[16] In the prompt limit ($\tau_N \equiv 0$), the approach taken by ATLAS would work. However, HNLs in the lowest two mass bins (5 and 10 GeV) have a small displacement, which can strongly affect the efficiency.

$ee\mu$ background			$\mu\mu e$ background		
b_i	b_i^-	b_i^+	b_i	b_i^-	b_i^+
19.0	14.9	23.1	21.3	17.5	25.1
18.0	14.4	21.7	13.8	10.6	17.0
21.0	17.4	24.7	18.7	15.1	22.3
13.6	10.9	16.2	13.3	10.6	16.1
6.1	4.2	7.8	13.1	10.2	15.9

Table 4.3.6: Background in 5 invariant mass bins (rows) for the searches in $e^\pm e^\pm \mu^\mp$ and $\mu^\pm \mu^\pm e^\mp$ channels correspondingly. The background is taken from Figure 5 in [86]. Only the total value (without individual contributions) is shown.

hypothesis. The full likelihood for the signal + background hypothesis is thus:

$$\mathcal{L}(x|H_{s+b}) = \mathcal{P}(\mu_{\text{tot}}|\mathcal{N}(1,\sigma_{\text{tot}})) \times \prod_{i=1}^{10} \mathcal{P}(x_i|\text{Pois}(\mu_{\text{tot}} b_i + s_i))$$
$$\text{where } \mu_{\text{tot}} = \frac{\sum_i (x_i - s_i)}{\sum_i b_i} \tag{4.23}$$

The likelihood for the background-only hypothesis H_b is obtained by setting the signal s_i to zero in eq. (4.23).

In order to validate our simplified statistical analysis, we can compare the limits that it produces to the limits obtained by ATLAS, when using the *exact same counts* as ATLAS (extracted again from figure 5 in ref. [86]). In order to perform this comparison, a few changes need to be made. First, we need to reintroduce the signal strength parameter μ. Second, we need to consider both channels separately. After making these changes, we obtain the limits shown in figure 4.3.4. The mean ratio between our limits and the ones from ATLAS is 64%, and the worst-case ratio is 42%. Although not fully satisfactory, this discrepancy should still be small enough to allow us to reliably compare limits which differ by an order of magnitude or more, as we will do in the next section. This is especially true when the reinterpreted limits are all computed using the same method.

4.4 Results

In this section, we present and analyze the reinterpreted exclusion limits obtained using the analysis described in section 4.3. We use the benchmark points defined in figure 4.2.1. These benchmarks have been chosen to represent both typical and extreme ratios of the mixing angles $U_e^2 : U_\mu^2 : U_\tau^2$.

The reinterpreted exclusion limits for the total mixing angle U_{tot}^2 (eq. (4.8)) are presented in figures 4.4.1 and 4.4.2, for the Majorana- and Dirac-like cases respectively. The limits for the individual mixing angles $U_\alpha^2 = \sum_{I=1,2} |\Theta_{\alpha I}|^2$ are presented in figures 4.4.3 to 4.4.6.

The legend, for all plots, is as follows. The thick dashed and dotted lines represent the exclusion limits obtained under the assumption of a Majorana-like HNL pair mixing with a single flavor. Up to a factor of 2, this corresponds

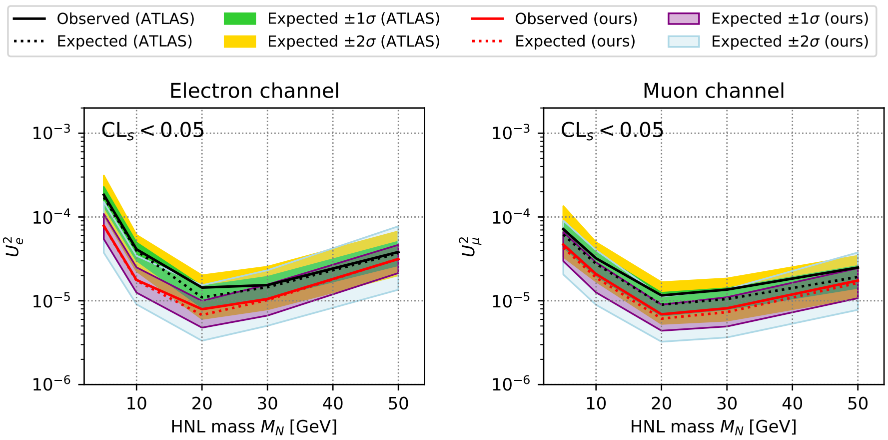

Figure 4.3.4: Comparison of the limits obtained using our simplified statistical model with the ones observed by ATLAS, using the exact same dataset (i.e. event counts, total background and expected signal).

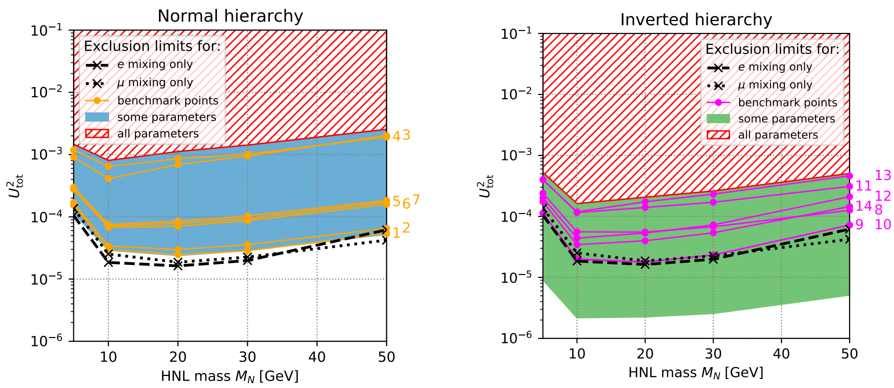

Figure 4.4.1: Original and reinterpreted exclusion limits (at 95% CL) on the total mixing angle $U^2_{\text{tot}} = \sum_{\alpha=e,\mu,\tau} \sum_{I=1,2} |\Theta_{\alpha I}|^2$ for a **Majorana-like** HNL pair, and for the normal (left) and inverted (right) hierarchies. The black lines denote the limits observed under the single-flavor assumption, while the solid colored lines denote those obtained for the benchmark points defined in figure 4.2.1. The filled area represents the set of all possible (benchmark-dependent) 95% exclusion limits when considering all the ratios of mixing angles allowed by neutrino oscillation data (corresponding to the 2σ region in figure 4.2.1). The red hatched area is excluded at CL > 95% for all possible ratios of mixing angles, and thus constitutes a benchmark-independent exclusion limit.

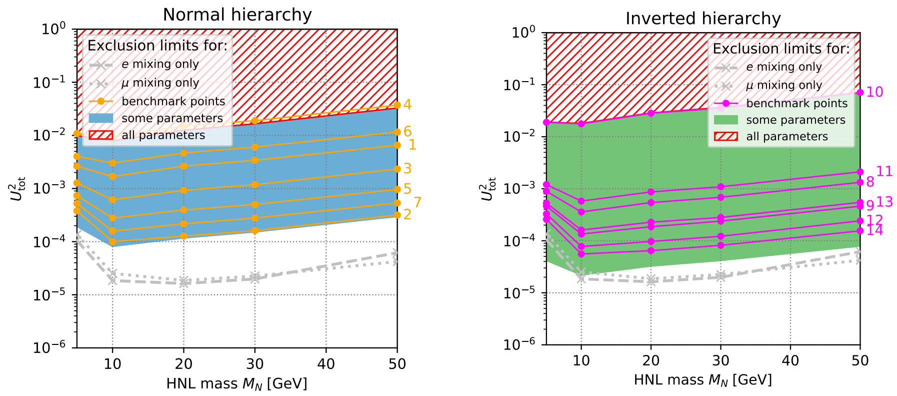

Figure 4.4.2: Same as figure 4.4.1, but for a **Dirac-like** HNL pair. The single-flavor mixing limits are grayed out because this search has *no sensitivity* to the Dirac-like case under this assumption; instead the limits for the Majorana-like case are given for comparison purpose.

to the scenario considered by ATLAS in the current prompt search. These limits are grayed out in the plots for the Dirac-like pair in order to emphasize that this search has *no sensitivity* to the Dirac-like case for single-flavor mixing. The solid colored lines denote the exclusion limits obtained for the various benchmark points defined in figure 4.2.1, and the benchmarks can be identified using the numbers in the right margin. All these limits correspond to the observed exclusion limit, and have all been derived using the same statistical method, which we described in 4.3.5. Therefore, they might slightly deviate from the actual ATLAS limits from ref. [86], but they should be comparable among themselves. The colored, filled area represents the set of possible (benchmark-dependent) limits spanned by all the combinations of mixing angles allowed by the NuFIT 5.0 neutrino data (at $2\sigma \approx 95\%$ CL[17]). In other words, it shows the dependence of the exclusion limits on the choice of the benchmark point, within the constraints from neutrino oscillation data which are represented by the similarly-colored area in figure 4.2.1. For each point within this area, there exists at least one valid combination of mixing angles which produces a limit at this level. Finally, the hatched red area denotes the set of mixing angles which are excluded for all allowed ratios of mixing angles. It thus represents the most conservative (benchmark-independent) limit that can be obtained for a given model[18]. No choice of mixing angles which is in agreement with neutrino oscillation data (within the 2 HNL seesaw model) can produce a limit within the red hatched region, for any of the mixing angles.

[17] The 2σ confidence limit is given assuming the specified hierarchy.

[18] The limits that we call "benchmark-dependent" rely on a specific set of model parameters (here: mixing angles), while the ones we call "benchmark-independent" have been obtained by marginalizing over these parameters; however, they still rely on the general properties of the model, such as the number of HNLs, the neutrino mass ordering or whether the HNLs behave as a Dirac-like or Majorana-like particle.

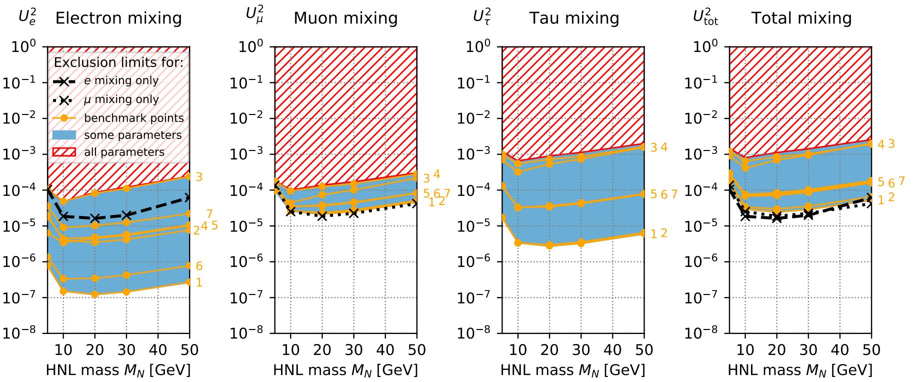

Figure 4.4.3: Original and reinterpreted exclusion limits (at $> 95\%$ CL) on the individual mixing angles $U_\alpha^2 = \sum_{I=1,2} |\Theta_{\alpha I}|^2 = x_\alpha U_{\rm tot}^2$ and the total mixing angle $U_{\rm tot}^2 = \sum_{\alpha=e,\mu,\tau} U_\alpha^2$ for a **Majorana-like** HNL pair and for the **normal hierarchy**. The legend is the same as in figure 4.4.1.

4.4.1 Majorana-like HNL pair

Let us first consider the case of a Majorana-like HNL pair, which is phenomenologically closest to the "single Majorana HNL" model considered by ATLAS and many other experiments. This corresponds to figures 4.4.1, 4.4.3 and 4.4.4. Apart from a trivial factor of two due to the fact that we have two nearly degenerate mass eigenstates, the only difference with ATLAS is that we consider a realistic seesaw model, which forces the HNLs to mix with all three mixing angles at the same time. Looking at the total mixing angle in figure 4.4.1, we immediately notice that the limits on $U_{\rm tot}^2$ are weaker than the single-flavor mixing limits for all our benchmarks, sometimes by more than an order of magnitude.[19] The pattern is obvious for the normal hierarchy (but also visible for the inverted one): the benchmark points which have the strongest tau fraction $x_\tau = U_\tau^2/U_{\rm tot}^2$ also have the worst sensitivity. This is the manifestation of a well-known phenomenon in non-minimal HNL models: the introduction of new decay channels (here mediated by the tau mixing) reduces the branching fraction of the HNLs into the search channels. This has an important consequence: exclusion limits derived for U_α^2 under the single-flavor assumption *do not translate directly* into limits on U_α^2 in a model where HNLs mix with multiple flavors.[20] Instead, such limits *must always be recast!*

When we look at the exclusion limits obtained for the individual mixing angles (in figure 4.4.3 for the normal hierarchy and figure 4.4.4 for the inverted hierarchy), we observe another striking effect: for specific benchmarks, the exclusion limits on individual mixing angles can sometimes be significantly *stronger* than the single-flavor limits. This actually reflects a rather trivial phenomenon: by fixing the $U_e^2 : U_\mu^2 : U_\tau^2$ ratio and setting a limit on one of the mixing angles (e.g. the largest, such as U_e^2 for benchmark 10 in the inverted

[19] It is also possible to obtain a stronger exclusion limit on $U_{\rm tot}^2$ (such as the lower limit of the green band in the inverted hierarchy), but this seems to occur only for very specific combinations of the mixing angles.

[20] Similarly, such limits do not apply if the HNLs have new interactions.

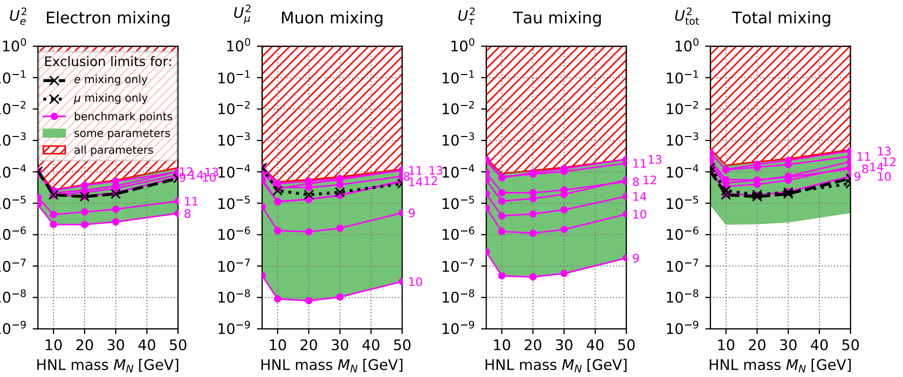

Figure 4.4.4: Same as figure 4.4.3, for a **Majorana-like** HNL pair and the **inverted hierarchy**.

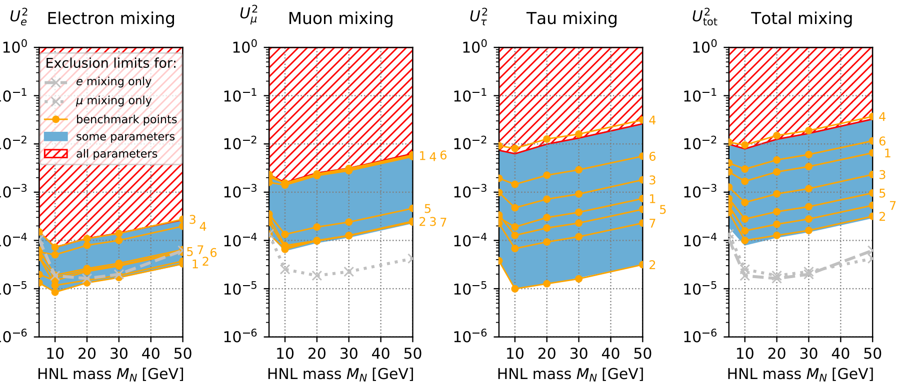

Figure 4.4.5: Same as figure 4.4.3, for a **Dirac-like** HNL pair and the **normal hierarchy**. The legend is the same as in figure 4.4.2.

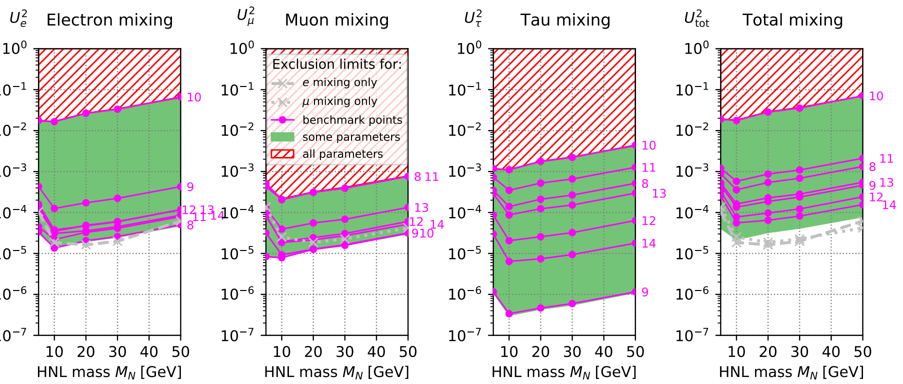

Figure 4.4.6: Same as figure 4.4.5, for a **Dirac-like** HNL pair and the **inverted hierarchy**.

hierarchy; see figure 4.4.4), we immediately obtain indirect limits on the other mixing angles (such as U_μ^2), which are enhanced by the ratio of the two mixing angles (in this case $U_\mu^2/U_e^2 = x_\mu/x_e \sim 1/2000$). Notice, in particular, how within this two-HNLs seesaw model we obtain an indirect limit (hatched region) on the *tau* mixing angle, which was not *directly* probed by this search. This simply reflects the fact that no valid combination of mixing angles which passes the constraints set by ATLAS in both the electron and muon channels, can have a mixing angle U_τ^2 with tau above this limit. Although the fact that introducing new constraints (such as fixing the ratio of mixing angles) can increase the sensitivity is not unexpected, it may still be useful when one considers *specific* sets of model parameters. This situation is not so far-fetched, since this is exactly what happens when performing a scan over the parameter space in order to e.g. combine constraints from multiple different sources. We would also expect future experimental results (such as excluding one neutrino mass hierarchy, or observing / setting limits on neutrinoless double-beta decay) to introduce additional constraints on the possible combinations of mixing angles, thus leading to a more predictive model. These potential use cases once again support the reinterpretation of exclusion limits.

4.4.2 Dirac-like HNL pair

Let us now turn our attention to the case of a Dirac-like HNL pair. Unlike in the Majorana-like case, there is *no* observable lepton number violation in this case, since the HNLs do not have enough time to oscillate among themselves. Its phenomenology thus significantly departs from the one of a single Majorana HNL, usually considered by experiments. In particular, the only lepton-number-conserving contributions to the experimental signatures considered in [86] come from processes in which the HNL mixes with different flavors during its production and decay (due to the veto of opposite-charge same-flavor trilepton events). This search therefore has *no sensitivity* to HNLs mixing with a single flavor!

By reinterpreting the limits (obtained for one Majorana HNL) within a realistic seesaw model (which *requires* HNLs to mix with all three flavors), we are nonetheless able to set some exclusion limits for this model. These limits are presented in figures 4.4.2, 4.4.5 and 4.4.6. The legend is the same as for the Majorana-like HNL pair, except for the single-flavor mixing limits which are grayed out in order to emphasize that they were computed for a different model (Majorana-like HNLs) and are only present here for comparison purpose. Looking at our benchmark points, we immediately notice that their limits for the total mixing angle (figure 4.4.2) are always weaker than the Majorana-like/single-flavor limits, sometimes by more than three orders of magnitude. The weakest limits are obtained when one of U_e^2 or U_μ^2 is suppressed compared to the other, which is unsurprising given that this approximates the single-flavor mixing case, to which we have no sensitivity. Looking at the filled area, we also observe a wider possible range of limits (with variations by more than two orders of magnitude) compared to the Majorana-like case, depending of the specific mixing pattern. This reflects the fact that the limits now depend mainly on two mixing angles instead of just one, which enhances the benchmark dependence. Finally, similarly to the Majorana-like case, we observe that we can obtain strong benchmark-dependent limits on the individual mixing angles (see figures 4.4.5 and 4.4.6), as well as some benchmark-independent limits (for this specific seesaw model with a Dirac-like HNL pair). The latter are significantly weaker (by up to two orders of magnitude) than for a Majorana-like HNL pair, due to the larger variation among benchmarks.

We can summarize the case of Dirac-like HNLs by emphasizing how, despite the absence of sensitivity to the single-flavor mixing case, we nonetheless managed to obtain both benchmark-dependent and benchmark-independent (but still model-dependent) exclusion limits by reinterpreting the ATLAS results within a realistic seesaw model featuring a Dirac-like HNL pair. Since the relevant processes now depend on the product of two different mixing angles, limits for Dirac-like HNLs are more benchmark dependent than for Majorana-like HNLs, resulting in weaker benchmark-independent exclusion limits (hatched area) for this model. Yet, the reinterpretation allowed us to obtain a limit on all three mixing angles (as well as their sum), where there was previously none.

4.5 Conclusion & outlook

4.5.1 Reinterpretation

Heavy neutral leptons (HNLs) are promising candidates for explaining neutrino masses and oscillations. Within the seesaw model, their mass scale is not predicted by neutrino masses. Experiments searching for HNLs typically report null results in the form of exclusion limits on the mixing angle with one of the lepton flavors. We emphasize that these constraints are neither model nor benchmark independent. Rather they correspond to limits obtained *within a specific model where one HNL mixes with a single flavor*. These simplified models are incompatible with the observed neutrino masses and mixing patterns. One may then wonder if the exclusion limits reported within these models remain valid when considering more realistic and theoretically motivated models of HNLs. In this work, we have performed a *reinterpretation* of the latest ATLAS prompt search for heavy neutral leptons [86] within one of the simplest *realistic* models: a low-scale seesaw mechanism with two quasi-degenerate HNLs. At least two HNLs are required in order to be compatible with neutrino oscillation data, and the combination of their mixing angles is constrained by the seesaw relation. In particular, for two HNLs, no mixing angle can be zero.

Our aim was to study to which extent the exclusion limits on the HNL mixing angles are model or benchmark dependent and by how much they change when considering our more realistic model.

Furthermore, the two HNLs must form a “quasi-Dirac” pair (i.e. be nearly degenerate, with a specific mixing pattern) for sufficiently large mixing angles (which may be accessible at current experiments) to be viable. Depending on the specific value of the mass splitting as well as the length scale over which the HNLs are observed, this quasi-Dirac pair may behave either as a Majorana-like or a Dirac-like particle, due to quantum interference between the two mass eigenstates. Only Majorana-like HNL pairs feature lepton number violating decays, and the different spin correlation patterns for LNC and LNV decay chains lead to different signal efficiencies for Majorana- and Dirac-like HNLs. Moreover, due to the veto applied by ATLAS on opposite-charge same-flavor lepton pairs, different diagrams, which depend on different combinations of mixing angles, contribute for Majorana- and Dirac-like HNLs. In particular, the only diagrams contributing to the signal in the case of Dirac-like HNLs involve two different mixing angles, such that there can be no sensitivity at all under the single-flavor mixing assumption! In order to handle both the Majorana- and Dirac-like cases, we perform the reinterpretation for each of them separately.

For Majorana-like HNL pairs, we have observed that:

- The exclusion limit on the total mixing angle U^2_{tot} is almost always weaker (sometimes by more that one order of magnitude) in realistic models than for single-flavor mixing. This is essentially caused by the opening of new decay channels (hence reducing the other branching fractions) which do not contribute to the search signature.
- Fixing the ratio of the mixing angles can result in (sometimes significantly) stronger *indirect* constraints on some of the mixing angles. This can be useful when performing scans over the models parameters.
- Assuming the two-HNLs seesaw model and marginalizing over the ratio of mixing angles while keeping the HNL mass fixed, we can obtain limits on the individual mixing angles (including the tau mixing angle, which was not probed directly by this search) which do not depend on the specific mixing pattern.

For Dirac-like HNLs pairs, we have observed that:

- Contrary to the single-flavor mixing where the signal was identically zero, in our realistic model no single mixing angle can ever be zero, which ensures that we can always set an indirect (model-dependent) limit.
- The limits on the total mixing angle are, however, almost always weaker (by up to three orders of magnitude) than in the Majorana-like/single-flavor case.
- The weakest limits are obtained when one of U^2_e or U^2_μ is suppressed compared to the other. This is expected, since these mixing patterns approximate the single-flavor case.
- Compared to the Majorana-like case, the dependence of the limits on the specific benchmark (mixing pattern) is stronger. This is likely caused by the fact that the product of two different mixing angles enters the cross section as a factor (instead of a single mixing angle) thus enhancing the parametric dependence.
- Similarly to the Majorana-like case, we can also set strong benchmark-dependent limits on the individual mixing angles by fixing their ratio. However, the corresponding marginalized/benchmark-independent limits are significantly weaker (by up to two orders of magnitude) due to the increased benchmark-dependence.

These results, despite having been obtained in the specific case of two heavy neutral leptons and the ATLAS experiment, emphasize the importance of always reinterpreting the experimental limits within the model of interest in order to obtain *reliable* exclusion limits. Failure to distinguish "limits on individual mixing angles obtained within a *realistic* model" from "limits obtained under the *assumption* of single-flavor mixing" can lead to incorrect conclusions! When assuming specific choices of model parameters (as in parameter scans), stronger constraints can be derived for the individual mixing angles. Benchmark-independent constraints can also be derived by marginalizing over the set of parameters allowed within a specific model. In order for experimental results to be truly usable for constraining models, experiments should strive to make their results readily reinterpretable within closely related models, keeping in mind that the main users of these results — theorists — are typically unfamiliar with the inner workings of the experiment. Below we outline a concrete proposal for reporting these results in the case of heavy neutral leptons, that would allow easily reinterpreting the exclusion limits. It is *not* precise enough to allow users to perform a *search*, for which only the full analysis should be used.

4.5.2 Recommendations for experiments

In order to produce reinterpretable exclusion limits on heavy neutral leptons, we recommend that experiments report the following data, ideally as data files[21] in a common format and with a documented structure left to their discretion:

- The observed bin counts.
- For the signal, both:
  1. The prompt efficiency $\epsilon_{0,P,b}(M_N) = \epsilon_{P,b}(M_N, \tau_N = 0)$ for every Feynman diagram P and every bin b in all signal regions. All possible processes which contribute to the search signature for *at least one* combination of the mixing angles should be included, in order to be able to compute the signal outside specific flavor mixing assumptions.
  2. If the parametrization in eq. (4.18) (or a modification thereof) allows reproducing the actual efficiency *even approximately*, report the relevant parameters such as the lifetime cutoff τ_0 (possibly for every process P if a single value does not give a good enough fit). The chosen parametrization should always have the same asymptotes as eq. (4.18).
- For the background, either:
  - The "full" likelihood, including every background component and nuisance parameter used in the analysis (to the extent that this is possible). This should ideally be reported as working code.
  - A simplified likelihood, containing only the dominant background components and nuisance parameters.
  - The covariance matrix for all the signal bins, across all signal regions.

  Regardless of their choice, we recommend that experiments validate the accuracy of the simplified background model by comparing the resulting limits with those obtained using the full analysis.

[21] As an example, the final publication will include the data files produced for this reinterpretation.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (GA 336581, 694896) and from the Carlsberg foundation. The work of IT has been supported by ERC-AdG-2015 grant 694896 and by the Swiss National Science Foundation Excellence grant 200020B_182864.

Bibliography

[1] P. Langacker, "Introduction to the Standard Model and Electroweak Physics," in *Theoretical Advanced Study Institute in Elementary Particle Physics: The Dawn of the LHC Era*, pp. 3–48. 2010. arXiv:0901.0241 [hep-ph].

[2] H. K. Dreiner, H. E. Haber, and S. P. Martin, "Two-component spinor techniques and Feynman rules for quantum field theory and supersymmetry," *Phys. Rept.* **494** (2010) 1–196, arXiv:0812.1594 [hep-ph].

[3] S. L. Glashow, "Partial Symmetries of Weak Interactions," *Nucl. Phys.* **22** (1961) 579–588.

[4] S. Weinberg, "A Model of Leptons," *Phys. Rev. Lett.* **19** (1967) 1264–1266.

[5] A. Salam, "Weak and Electromagnetic Interactions," *Conf. Proc. C* **680519** (1968) 367–377.

[6] S. L. Glashow, J. Iliopoulos, and L. Maiani, "Weak Interactions with Lepton-Hadron Symmetry," *Phys. Rev. D* **2** (1970) 1285–1292.

[7] M. Srednicki, *Quantum field theory*. Cambridge University Press, 1, 2007.

[8] F. Englert and R. Brout, "Broken Symmetry and the Mass of Gauge Vector Mesons," *Phys. Rev. Lett.* **13** (1964) 321–323.

[9] P. W. Higgs, "Broken Symmetries and the Masses of Gauge Bosons," *Phys. Rev. Lett.* **13** (1964) 508–509.

[10] P. W. Higgs, "Spontaneous Symmetry Breakdown without Massless Bosons," *Phys. Rev.* **145** (1966) 1156–1163.

[11] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, "Global Conservation Laws and Massless Particles," *Phys. Rev. Lett.* **13** (1964) 585–587.

[12] **ATLAS** Collaboration, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC," *Phys. Lett. B* **716** (2012) 1–29, arXiv:1207.7214 [hep-ex].

[13] **CMS** Collaboration, "Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC," *Phys. Lett. B* **716** (2012) 30–61, arXiv:1207.7235 [hep-ex].

[14] R. Davis, Jr., D. S. Harmer, and K. C. Hoffman, "Search for neutrinos from the sun," *Phys. Rev. Lett.* **20** (1968) 1205–1209.

[15] **Super-Kamiokande** Collaboration, Y. Fukuda *et al.*, "Evidence for oscillation of atmospheric neutrinos," *Phys. Rev. Lett.* **81** (1998) 1562–1567, arXiv:hep-ex/9807003.

[16] **SNO** Collaboration, "Measurement of the rate of $\nu_e + d \to p + p + e^-$ interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory," *Phys. Rev. Lett.* **87** (2001) 071301, arXiv:nucl-ex/0106015.

[17] B. Pontecorvo, "Mesonium and Antimesonium," *Soviet Journal of Experimental and Theoretical Physics* **6** (1958) 429.

[18] B. Pontecorvo, "Neutrino Experiments and the Problem of Conservation of Leptonic Charge," *Soviet Journal of Experimental and Theoretical Physics* **26** (May, 1968) 984.

[19] L. Wolfenstein, "Neutrino Oscillations in Matter," *Phys. Rev. D* **17** (1978) 2369–2374.

[20] S. P. Mikheyev and A. Y. Smirnov, "Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos," *Sov. J. Nucl. Phys.* **42** (1985) 913–917.

[21] Z. Maki, M. Nakagawa, and S. Sakata, "Remarks on the unified model of elementary particles," *Prog. Theor. Phys.* **28** (1962) 870–880.

[22] M. Beuthe, "Oscillations of neutrinos and mesons in quantum field theory," *Phys. Rept.* **375** (2003) 105–218, arXiv:hep-ph/0109119 [hep-ph].

[23] L. Canetti, M. Drewes, and M. Shaposhnikov, "Matter and Antimatter in the Universe," *New J. Phys.* **14** (2012) 095012, arXiv:1204.4186 [hep-ph].

[24] A. D. Sakharov, "Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe," *Sov. Phys. Usp.* **34** no. 5, (1991) 392–393.

[25] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov, "On the Anomalous Electroweak Baryon Number Nonconservation in the Early Universe," *Phys. Lett. B* **155** (1985) 36.

[26] **Particle Data Group** Collaboration, P. Zyla *et al.*, "Review of Particle Physics," *PTEP* **2020** no. 8, (2020) 083C01.

[27] P. Minkowski, "$\mu \to e\gamma$ at a Rate of One Out of 10^9 Muon Decays?," *Phys. Lett.* **67B** (1977) 421–428.

[28] T. Yanagida, "Horizontal Gauge Symmetry and Masses of Neutrinos," *Conf. Proc.* **C7902131** (1979) 95–99.

[29] S. L. Glashow, "The Future of Elementary Particle Physics," *NATO Sci. Ser. B* **61** (1980) 687.

[30] M. Gell-Mann, P. Ramond, and R. Slansky, "Complex Spinors and Unified Theories," *Conf. Proc.* **C790927** (1979) 315–321, `arXiv:1306.4669 [hep-th]`.

[31] R. N. Mohapatra and G. Senjanovic, "Neutrino Mass and Spontaneous Parity Nonconservation," *Phys. Rev. Lett.* **44** (1980) 912. [,231(1979)].

[32] R. N. Mohapatra and G. Senjanovic, "Neutrino Masses and Mixings in Gauge Models with Spontaneous Parity Violation," *Phys. Rev.* **D23** (1981) 165.

[33] J. Schechter and J. W. F. Valle, "Neutrino Masses in $SU(2) \times \mathrm{U}(1)$ Theories," *Phys. Rev.* **D22** (1980) 2227.

[34] J. Schechter and J. W. F. Valle, "Neutrino Decay and Spontaneous Violation of Lepton Number," *Phys. Rev.* **D25** (1982) 774.

[35] T. Asaka and M. Shaposhnikov, "The νMSM, Dark Matter and Baryon Asymmetry of the Universe," *Phys. Lett.* **B620** (2005) 17–26, `arXiv:hep-ph/0505013 [hep-ph]`.

[36] T. Asaka, S. Blanchet, and M. Shaposhnikov, "The νMSM, Dark Matter and Neutrino Masses," *Phys. Lett.* **B631** (2005) 151–156, `arXiv:hep-ph/0503065 [hep-ph]`.

[37] S. Eijima, M. Shaposhnikov, and I. Timiryasov, "Parameter space of baryogenesis in the νMSM," *JHEP* **07** (2019) 077, `arXiv:1808.10833 [hep-ph]`.

[38] M. Shaposhnikov, "The νMSM, leptonic asymmetries, and properties of singlet fermions," *JHEP* **08** (2008) 008, `arXiv:0804.4542 [hep-ph]`.

[39] **SHiP** Collaboration, "Sensitivity of the SHiP experiment to Heavy Neutral Leptons," *JHEP* **04** (2019) 077, `arXiv:1811.00930 [hep-ph]`.

[40] **FCC-ee study Team** Collaboration, A. Blondel, E. Graverini, N. Serra, and M. Shaposhnikov, "Search for Heavy Right Handed Neutrinos at the FCC-ee," *Nucl. Part. Phys. Proc.* **273-275** (2016) 1883–1890, `arXiv:1411.5230 [hep-ex]`.

[41] D. Curtin *et al.*, "Long-Lived Particles at the Energy Frontier: The MATHUSLA Physics Case," *Rept. Prog. Phys.* **82** no. 11, (2019) 116201, `arXiv:1806.07396 [hep-ph]`.

[42] M. Drewes, J. Hajer, J. Klaric, and G. Lanfranchi, "NA62 sensitivity to heavy neutral leptons in the low scale seesaw model," *JHEP* **07** (2018) 105, `arXiv:1801.04207 [hep-ph]`.

[43] D. Dercks, H. K. Dreiner, M. Hirsch, and Z. S. Wang, "Long-Lived Fermions at AL3X," *Phys. Rev.* **D99** no. 5, (2019) 055020, arXiv:1811.01995 [hep-ph].

[44] **FASER** Collaboration, "FASER's physics reach for long-lived particles," *Phys. Rev. D* **99** no. 9, (2019) 095011, arXiv:1811.12522 [hep-ph].

[45] E. K. Akhmedov, V. A. Rubakov, and A. Yu. Smirnov, "Baryogenesis via neutrino oscillations," *Phys. Rev. Lett.* **81** (1998) 1359–1362, arXiv:hep-ph/9803255 [hep-ph].

[46] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*. 2007.

[47] T. Takagi *Japanese J. Math.* **1** (1927) 83.

[48] J. Kersten and A. Yu. Smirnov, "Right-Handed Neutrinos at CERN LHC and the Mechanism of Neutrino Mass Generation," *Phys. Rev.* **D76** (2007) 073005, arXiv:0705.3221 [hep-ph].

[49] A. Boyarsky, O. Ruchayskiy, and M. Shaposhnikov, "The Role of sterile neutrinos in cosmology and astrophysics," *Ann. Rev. Nucl. Part. Sci.* **59** (2009) 191–214, arXiv:0901.0011 [hep-ph].

[50] A. Boyarsky, M. Drewes, T. Lasserre, S. Mertens, and O. Ruchayskiy, "Sterile Neutrino Dark Matter," *Prog. Part. Nucl. Phys.* **104** (2019) 1–45, arXiv:1807.07938 [hep-ph].

[51] J. A. Casas and A. Ibarra, "Oscillating neutrinos and $\mu \to e, \gamma$," *Nucl. Phys.* **B618** (2001) 171–204, arXiv:hep-ph/0103065 [hep-ph].

[52] Z.-z. Xing, "Casas-Ibarra Parametrization and Unflavored Leptogenesis," *Chin. Phys. C* **34** (2010) 1–6, arXiv:0902.2469 [hep-ph].

[53] A. Abada, S. Davidson, A. Ibarra, F.-X. Josse-Michaux, M. Losada, and A. Riotto, "Flavour Matters in Leptogenesis," *JHEP* **09** (2006) 010, arXiv:hep-ph/0605281.

[54] 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," *JHEP* **09** (2020) 178, arXiv:2007.14792 [hep-ph].

[55] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou, "NuFIT 5.0 (2020)," 2020. http://www.nu-fit.org.

[56] J.-L. Tastet and I. Timiryasov, "Dirac vs. Majorana HNLs (and their oscillations) at SHiP," *JHEP* **04** (2020) 005, arXiv:1912.05520 [hep-ph].

[57] R. Ruiz, "Quantitative study on helicity inversion in Majorana neutrino decays at the LHC," *Phys. Rev. D* **103** no. 1, (2021) 015022, arXiv:2008.01092 [hep-ph].

[58] E. Majorana, "Teoria simmetrica dell'elettrone e del positrone," *Nuovo Cim.* **14** (1937) 171–184.

[59] S. M. Bilenky and B. Pontecorvo, "Again on Neutrino Oscillations," *Lett. Nuovo Cim.* **17** (1976) 569.

[60] S. M. Bilenky and B. Pontecorvo, "Lepton Mixing and Neutrino Oscillations," *Phys. Rept.* **41** (1978) 225–261.

[61] Y. F. Perez and C. J. Quimbay, "Majorana Neutrino Oscillations in Vacuum," *J. Mod. Phys.* **3** (2012) 803–814, `arXiv:1103.2781 [hep-ph]`.

[62] E. Akhmedov, *Majorana neutrinos and other Majorana particles: Theory and experiment.* 12, 2014. `arXiv:1412.3320 [hep-ph]`.

[63] D. Boyanovsky, "Nearly degenerate heavy sterile neutrinos in cascade decay: mixing and oscillations," *Phys. Rev.* **D90** no. 10, (2014) 105024, `arXiv:1409.4265 [hep-ph]`.

[64] G. Cvetic, C. S. Kim, R. Kogerler, and J. Zamora-Saa, "Oscillation of heavy sterile neutrino in decay of $B \to \mu e \pi$," *Phys. Rev.* **D92** (2015) 013015, `arXiv:1505.04749 [hep-ph]`.

[65] G. Anamiati, M. Hirsch, and E. Nardi, "Quasi-Dirac neutrinos at the LHC," *JHEP* **10** (2016) 010, `arXiv:1607.05641 [hep-ph]`.

[66] S. Antusch, E. Cazzato, and O. Fischer, "Resolvable heavy neutrino–antineutrino oscillations at colliders," *Mod. Phys. Lett. A* **34** no. 07n08, (2019) 1950061, `arXiv:1709.03797 [hep-ph]`.

[67] J.-L. Tastet, "Sterile neutrino oscillations at the SHiP experiment," Master's thesis, École Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland, 2017.

[68] P. Hernández, J. Jones-Pérez, and O. Suarez-Navarro, "Majorana vs Pseudo-Dirac Neutrinos at the ILC," *Eur. Phys. J. C* **79** no. 3, (2019) 220, `arXiv:1810.07210 [hep-ph]`.

[69] G. Cvetič, A. Das, and J. Zamora-Saá, "Probing heavy neutrino oscillations in rare W boson decays," *J. Phys. G* **46** (2019) 075002, `arXiv:1805.00070 [hep-ph]`.

[70] P. S. Bhupal Dev, R. N. Mohapatra, and Y. Zhang, "CP Violating Effects in Heavy Neutrino Oscillations: Implications for Colliders and Leptogenesis," *JHEP* **11** (2019) 137, `arXiv:1904.04787 [hep-ph]`.

[71] G. Cvetič, A. Das, S. Tapia, and J. Zamora-Saá, "Measuring the heavy neutrino oscillations in rare W boson decays at the Large Hadron Collider," *J. Phys.* **G47** no. 1, (2020) 015001, `arXiv:1905.03097 [hep-ph]`.

[72] S. Tapia and J. Zamora-Saá, "Exploring CP-Violating heavy neutrino oscillations in rare tau decays at Belle II," *Nucl. Phys. B* **952** (2020) 114936, `arXiv:1906.09470 [hep-ph]`.

[73] R. Sachs, "Interference phenomena of neutral K mesons," *Annals of Physics* **22** no. 2, (1963) 239 – 262. http://www.sciencedirect.com/science/article/pii/0003491663900551.

[74] C. Giunti, C. W. Kim, J. A. Lee, and U. W. Lee, "On the treatment of neutrino oscillations without resort to weak eigenstates," *Phys. Rev.* **D48** (1993) 4310–4317, arXiv:hep-ph/9305276 [hep-ph].

[75] E. K. Akhmedov and J. Kopp, "Neutrino Oscillations: Quantum Mechanics vs. Quantum Field Theory," *JHEP* **04** (2010) 008, arXiv:1001.4815 [hep-ph]. [Erratum: JHEP10,052(2013)].

[76] E. K. Akhmedov and A. Yu. Smirnov, "Neutrino oscillations: Entanglement, energy-momentum conservation and QFT," *Found. Phys.* **41** (2011) 1279–1306, arXiv:1008.2077 [hep-ph].

[77] E. Akhmedov, "Quantum mechanics aspects and subtleties of neutrino oscillations," in *International Conference on History of the Neutrino: 1930-2018 Paris, France, September 5-7, 2018*. 2019. arXiv:1901.05232 [hep-ph].

[78] A. G. Cohen, S. L. Glashow, and Z. Ligeti, "Disentangling Neutrino Oscillations," *Phys. Lett. B* **678** (2009) 191–196, arXiv:0810.4602 [hep-ph].

[79] S. Antusch and J. Rosskopp, "Heavy Neutrino-Antineutrino Oscillations in Quantum Field Theory," arXiv:2012.05763 [hep-ph].

[80] A. Das, P. S. B. Dev, and R. N. Mohapatra, "Same Sign versus Opposite Sign Dileptons as a Probe of Low Scale Seesaw Mechanisms," *Phys. Rev.* **D97** no. 1, (2018) 015018, arXiv:1709.06553 [hep-ph].

[81] F. Najafi, J. Kumar, and D. London, "CP Violation in Rare Lepton-Number-Violating W Decays at the LHC," arXiv:2011.03686 [hep-ph].

[82] L. Canetti, M. Drewes, T. Frossard, and M. Shaposhnikov, "Dark Matter, Baryogenesis and Neutrino Oscillations from Right Handed Neutrinos," *Phys. Rev.* **D87** (2013) 093006, arXiv:1208.4607 [hep-ph].

[83] B. Kayser, "On the Quantum Mechanics of Neutrino Oscillation," *Phys. Rev. D* **24** (1981) 110.

[84] K. Kiers and N. Weiss, "Neutrino oscillations in a model with a source and detector," *Phys. Rev. D* **57** (1998) 3091–3105, arXiv:hep-ph/9710289.

[85] K. Bondarenko, A. Boyarsky, J. Klaric, O. Mikulenko, O. Ruchayskiy, V. Syvolap, and I. Timiryasov, "An allowed window for heavy neutral leptons below the kaon mass," arXiv:2101.09255 [hep-ph].

[86] **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," *JHEP* **10** (2019) 265, arXiv:1905.09787 [hep-ex].

[87] F. Thiele, "An ATLAS Search for Sterile Neutrinos," May, 2019. https://cds.cern.ch/record/2676324.

[88] T. Yanagida, "Horizontal Symmetry and Masses of Neutrinos," *Prog. Theor. Phys.* **64** (1980) 1103.

[89] F. Vissani, "Do experiments suggest a hierarchy problem?," *Phys. Rev.* **D57** (1998) 7027–7030, arXiv:hep-ph/9709409 [hep-ph].

[90] F. Bezrukov, M. Yu. Kalmykov, B. A. Kniehl, and M. Shaposhnikov, "Higgs Boson Mass and New Physics," *JHEP* **10** (2012) 140, arXiv:1205.2893 [hep-ph]. [,275(2012)].

[91] J. Beacham *et al.*, "Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report," *J. Phys.* **G47** no. 1, (2020) 010501, arXiv:1901.09966 [hep-ex].

[92] J. P. Chou, D. Curtin, and H. J. Lubatti, "New Detectors to Explore the Lifetime Frontier," *Phys. Lett.* **B767** (2017) 29–36, arXiv:1606.06298 [hep-ph].

[93] **MATHUSLA** Collaboration, "A Letter of Intent for MATHUSLA: A Dedicated Displaced Vertex Detector above ATLAS or CMS," arXiv:1811.00927 [physics.ins-det].

[94] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, "ForwArd Search ExpeRiment at the LHC," *Phys. Rev.* **D97** no. 3, (2018) 035001, arXiv:1708.09389 [hep-ph].

[95] F. Kling and S. Trojanowski, "Heavy Neutral Leptons at FASER," *Phys. Rev.* **D97** no. 9, (2018) 095016, arXiv:1801.08947 [hep-ph].

[96] **FASER** Collaboration, "Letter of Intent for FASER: ForwArd Search ExpeRiment at the LHC," arXiv:1811.10243 [physics.ins-det].

[97] V. V. Gligorov, S. Knapen, M. Papucci, and D. J. Robinson, "Searching for Long-lived Particles: A Compact Detector for Exotics at LHCb," *Phys. Rev.* **D97** no. 1, (2018) 015023, arXiv:1708.09395 [hep-ph].

[98] G. Aielli *et al.*, "Expression of interest for the CODEX-b detector," *Eur. Phys. J. C* **80** no. 12, (2020) 1177, arXiv:1911.00481 [hep-ex].

[99] V. V. Gligorov, S. Knapen, B. Nachman, M. Papucci, and D. J. Robinson, "Leveraging the ALICE/L3 cavern for long-lived particle searches," *Phys. Rev.* **D99** no. 1, (2019) 015023, arXiv:1810.03636 [hep-ph].

[100] M. Bauer, O. Brandt, L. Lee, and C. Ohm, "ANUBIS: Proposal to search for long-lived neutral particles in CERN service shafts," arXiv:1909.13022 [physics.ins-det].

[101] **LBNE** Collaboration, "The 2010 Interim Report of the Long-Baseline Neutrino Experiment Collaboration Physics Working Groups," arXiv:1110.6249 [hep-ex].

[102] I. Krasnov, "DUNE prospects in the search for sterile neutrinos," *Phys. Rev.* **D100** no. 7, (2019) 075023, arXiv:1902.06099 [hep-ph].

[103] P. Ballett, T. Boschi, and S. Pascoli, "Heavy Neutral Leptons from low-scale seesaws at the DUNE Near Detector," *JHEP* **03** (2020) 111, arXiv:1905.00284 [hep-ph].

[104] **NA62** Collaboration, "The Beam and detector of the NA62 experiment at CERN," *JINST* **12** no. 05, (2017) P05025, arXiv:1703.08501 [physics.ins-det].

[105] **SHiP** Collaboration, "A facility to Search for Hidden Particles (SHiP) at the CERN SPS," arXiv:1504.04956 [physics.ins-det].

[106] S. Alekhin *et al.*, "A facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case," *Rept. Prog. Phys.* **79** no. 12, (2016) 124201, arXiv:1504.04855 [hep-ph].

[107] J. C. Helo, M. Hirsch, and Z. S. Wang, "Heavy neutral fermions at the high-luminosity LHC," *JHEP* **07** (2018) 056, arXiv:1803.02212 [hep-ph].

[108] I. Boiarska, K. Bondarenko, A. Boyarsky, S. Eijima, M. Ovchynnikov, O. Ruchayskiy, and I. Timiryasov, "Probing baryon asymmetry of the Universe at LHC and SHiP," arXiv:1902.04535 [hep-ph].

[109] K. Bondarenko, A. Boyarsky, M. Ovchynnikov, and O. Ruchayskiy, "Sensitivity of the intensity frontier experiments for neutrino and scalar portals: analytic estimates," *JHEP* **08** (2019) 061, arXiv:1902.06240 [hep-ph].

[110] E. J. Chun, A. Das, S. Mandal, M. Mitra, and N. Sinha, "Sensitivity of Lepton Number Violating Meson Decays in Different Experiments," *Phys. Rev.* **D100** no. 9, (2019) 095022, arXiv:1908.09562 [hep-ph].

[111] M. Shaposhnikov, "A Possible Symmetry of the νMSM," *Nucl. Phys.* **B763** (2007) 49–59, arXiv:hep-ph/0605047 [hep-ph].

[112] K. Moffat, S. Pascoli, and C. Weiland, "Equivalence between massless neutrinos and lepton number conservation in fermionic singlet extensions of the Standard Model," arXiv:1712.07611 [hep-ph].

[113] R. N. Mohapatra, "Mechanism for Understanding Small Neutrino Mass in Superstring Theories," *Phys. Rev. Lett.* **56** (1986) 561–563.

[114] R. N. Mohapatra and J. W. F. Valle, "Neutrino Mass and Baryon Number Nonconservation in Superstring Models," *Phys. Rev.* **D34** (1986) 1642. [,235(1986)].

[115] E. K. Akhmedov, M. Lindner, E. Schnapka, and J. W. F. Valle, "Left-right symmetry breaking in NJL approach," *Phys. Lett.* **B368** (1996) 270–280, arXiv:hep-ph/9507275 [hep-ph].

[116] E. K. Akhmedov, M. Lindner, E. Schnapka, and J. W. F. Valle, "Dynamical left-right symmetry breaking," *Phys. Rev.* **D53** (1996) 2752–2780, arXiv:hep-ph/9509255 [hep-ph].

[117] M. Drewes, J. Klarić, and P. Klose, "On Lepton Number Violation in Heavy Neutrino Decays at Colliders," *JHEP* **11** (2019) 032, arXiv:1907.13034 [hep-ph]. [JHEP19,032(2020)].

[118] C. Arbelaéz, C. Dib, I. Schmidt, and J. C. Vasquez, "Probing the Dirac or Majorana nature of the Heavy Neutrinos in pure leptonic decays at the LHC," *Phys. Rev.* **D97** no. 5, (2018) 055011, arXiv:1712.08704 [hep-ph].

[119] C. O. Dib, C. S. Kim, and K. Wang, "Signatures of Dirac and Majorana sterile neutrinos in trilepton events at the LHC," *Phys. Rev.* **D95** no. 11, (2017) 115020, arXiv:1703.01934 [hep-ph].

[120] G. Cvetic, C. Dib, and C. S. Kim, "Probing Majorana neutrinos in rare $\pi^+ \to e^+ e^+ \mu^- \nu$ decays," *JHEP* **06** (2012) 149, arXiv:1203.0573 [hep-ph].

[121] A. B. Balantekin, A. de Gouvêa, and B. Kayser, "Addressing the Majorana vs. Dirac Question with Neutrino Decays," *Phys. Lett.* **B789** (2019) 488–495, arXiv:1808.10518 [hep-ph].

[122] A. Abada, C. Hati, X. Marcano, and A. M. Teixeira, "Interference effects in LNV and LFV semileptonic decays: the Majorana hypothesis," *JHEP* **09** (2019) 017, arXiv:1904.05367 [hep-ph].

[123] **LHCb** Collaboration, "Search for Majorana neutrinos in $B^- \to \pi^+ \mu^- \mu^-$ decays," *Phys. Rev. Lett.* **112** no. 13, (2014) 131802, arXiv:1401.5361 [hep-ex].

[124] **ATLAS** Collaboration, "Search for heavy Majorana neutrinos with the ATLAS detector in pp collisions at $\sqrt{s} = 8$ TeV," *JHEP* **07** (2015) 162, arXiv:1506.06020 [hep-ex].

[125] **CMS** Collaboration, V. Khachatryan *et al.*, "Search for heavy Majorana neutrinos in $\mu^\pm \mu^\pm +$ jets events in proton-proton collisions at $\sqrt{s} = 8$ TeV," *Phys. Lett.* **B748** (2015) 144–166, arXiv:1501.05566 [hep-ex].

[126] **CMS** Collaboration, "Search for heavy Majorana neutrinos in $\mathrm{e}^\pm \mathrm{e}^\pm +$ jets and $\mathrm{e}^\pm$ $\mu^\pm +$ jets events in proton-proton collisions at $\sqrt{s} = 8$ TeV," *JHEP* **04** (2016) 169, arXiv:1603.02248 [hep-ex].

[127] **LEBC-EHS** Collaboration, "*D* Meson Production From 400 GeV/*c* *pp* Interactions," *Phys. Lett.* **B189** (1987) 476. [Erratum: Phys. Lett.B208,530(1988)].

[128] **SHiP** Collaboration, "Measurement of associated charm production induced by 400 GeV/c protons," Tech. Rep. CERN-SPSC-2017-033. SPSC-EOI-017, CERN, Geneva, Oct, 2017. https://cds.cern.ch/record/2286844.

[129] K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy, "Phenomenology of GeV-scale Heavy Neutral Leptons," *JHEP* **11** (2018) 032, arXiv:1805.08567 [hep-ph].

[130] **SHiP** Collaboration, B. Hosseini and W. M. Bonivento, "Particle Identification tools and performance in the SHiP Experiment," Jul, 2017. https://cds.cern.ch/record/2282039.

[131] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "Lightgbm: A highly efficient gradient boosting decision tree," in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds., pp. 3146–3154. Curran Associates, Inc., 2017. http://papers.nips.cc/paper/6907-lightgbm-a-highly-efficient-gradient-boosting-decision-tree.pdf.

[132] Microsoft Corporation, "LightGBM," *GitHub repository* (2016–2019) . https://github.com/microsoft/LightGBM.

[133] J. Neyman and E. S. Pearson, "On the Problem of the Most Efficient Tests of Statistical Hypotheses," *Phil. Trans. Roy. Soc. Lond.* **A231** no. 694-706, (1933) 289–337.

[134] S. S. Wilks, "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses," *Annals Math. Statist.* **9** no. 1, (1938) 60–62.

[135] M. Drewes, "On the Minimal Mixing of Heavy Neutrinos," arXiv:1904.11959 [hep-ph].

[136] F. James, "Monte-Carlo phase space," 1968. https://cds.cern.ch/record/275743.

[137] P. Ilten, *Electroweak and Higgs Measurements Using Tau Final States with the LHCb Detector*. PhD thesis, University Coll., Dublin, 2013-09-06. arXiv:1401.4902 [hep-ex].

[138] S. Algeri, J. Aalbers, K. Dundas Morå, and J. Conrad, "Searching for new physics with profile likelihoods: Wilks and beyond," arXiv:1911.10237 [physics.data-an].

[139] P. Peebles, "Growth of the nonbaryonic dark matter theory," *Nature Astron.* **1** no. 3, (2017) 0057, arXiv:1701.05837 [astro-ph.CO].

[140] S. Bilenky, "Neutrino in Standard Model and beyond," *Phys. Part. Nucl.* **46** no. 4, (2015) 475–496, arXiv:1501.00232 [hep-ph].

[141] M. Shaposhnikov, "Is there a new physics between electroweak and Planck scales?," in *Astroparticle Physics: Current Issues, 2007 (APCI07)*. 8, 2007. arXiv:0708.3550 [hep-th].

[142] R. K. Ellis *et al.*, "Physics Briefing Book: Input for the European Strategy for Particle Physics Update 2020," arXiv:1910.11775 [hep-ex].

[143] C. Giganti, S. Lavignac, and M. Zito, "Neutrino oscillations: the rise of the PMNS paradigm," *Prog. Part. Nucl. Phys.* **98** (2018) 1–54, arXiv:1710.00715 [hep-ex].

[144] R. Foot, H. Lew, X. He, and G. C. Joshi, "Seesaw Neutrino Masses Induced by a Triplet of Leptons," *Z. Phys. C* **44** (1989) 441.

[145] A. Pilaftsis, "Radiatively induced neutrino masses and large Higgs neutrino couplings in the standard model with Majorana fields," *Z. Phys. C* **55** (1992) 275–282, arXiv:hep-ph/9901206.

[146] E. Accomando, L. Delle Rose, S. Moretti, E. Olaiya, and C. H. Shepherd-Themistocleous, "Novel SM-like Higgs decay into displaced heavy neutrino pairs in $U(1)'$ models," *JHEP* **04** (2017) 081, arXiv:1612.05977 [hep-ph].

[147] A. Caputo, P. Hernandez, J. Lopez-Pavon, and J. Salvado, "The seesaw portal in testable models of neutrino masses," *JHEP* **06** (2017) 112, arXiv:1704.08721 [hep-ph].

[148] L. Basso, A. Belyaev, S. Moretti, and C. H. Shepherd-Themistocleous, "Phenomenology of the minimal B-L extension of the Standard model: Z' and neutrinos," *Phys. Rev.* **D80** (2009) 055030, arXiv:0812.4313 [hep-ph].

[149] S. N. Gninenko, "A resolution of puzzles from the LSND, KARMEN, and MiniBooNE experiments," *Phys. Rev. D* **83** (2011) 015015, arXiv:1009.5536 [hep-ph].

[150] B. Batell, M. Pospelov, and B. Shuve, "Shedding Light on Neutrino Masses with Dark Forces," *JHEP* **08** (2016) 052, arXiv:1604.06099 [hep-ph].

[151] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, "Dipole Portal to Heavy Neutral Leptons," *Phys. Rev. D* **98** no. 11, (2018) 115015, arXiv:1803.03262 [hep-ph].

[152] O. Fischer, A. Hernández-Cabezudo, and T. Schwetz, "Explaining the MiniBooNE excess by a decaying sterile neutrino with mass in the 250 MeV range," *Phys. Rev. D* **101** no. 7, (2020) 075045, arXiv:1909.09561 [hep-ph].

[153] M. Chala and A. Titov, "One-loop matching in the SMEFT extended with a sterile neutrino," *JHEP* **05** (2020) 139, arXiv:2001.07732 [hep-ph].

[154] **NA62** Collaboration, "Search for production of an invisible dark photon in π^0 decays," *JHEP* **05** (2019) 182, arXiv:1903.08767 [hep-ex].

[155] **NA62** Collaboration, "Search for heavy neutral lepton production in K^+ decays to positrons," *Phys. Lett. B* **807** (2020) 135599, arXiv:2005.09575 [hep-ex].

[156] **NA62** Collaboration, B. Döbrich, "Searches for very weakly-coupled particles beyond the Standard Model with NA62," in *Proceedings of the 13th Patras Workshop on Axions, WIMPs and WISPs, PATRAS 2017*. 2017. arXiv:1711.08967 [hep-ex].

[157] **NA62** Collaboration, G. Lanfranchi, "Search for Hidden Sector particles at NA62," *PoS* **EPS-HEP2017** (2017) 301.

[158] **SHiP** Collaboration, P. Mermod, "Hidden sector searches with SHiP and NA62," in *2017 International Workshop on Neutrinos from Accelerators (NuFact17) Uppsala University Main Building, Uppsala, Sweden, September 25-30, 2017*. 2017. arXiv:1712.01768 [hep-ex].

[159] S. Gori, G. Perez, and K. Tobioka, "KOTO vs. NA62 Dark Scalar Searches," *JHEP* **08** (2020) 110, arXiv:2005.05170 [hep-ph].

[160] **PIENU** Collaboration, "Improved search for heavy neutrinos in the decay $\pi \to e\nu$," *Phys. Rev. D* **97** no. 7, (2018) 072012, arXiv:1712.03275 [hep-ex].

[161] G. Bernardi, G. Carugno, J. Chauveau, F. Dicarlo, M. Dris, J. Dumarchez, M. Ferro-Luzzi, J.-M. Levy, D. Lukas, J.-M. Perreau, Y. Pons, A.-M. Touchard, and F. Vannucci, "Further limits on heavy neutrino couplings," *Physics Letters B* **203** no. 3, (1988) 332 – 334.

[162] O. Ruchayskiy and A. Ivashko, "Experimental bounds on sterile neutrino mixing angles," *JHEP* **06** (2012) 100, arXiv:1112.3319 [hep-ph].

[163] D. Gorbunov and M. Shaposhnikov, "How to find neutral leptons of the νMSM?," *JHEP* **10** (2007) 015, arXiv:0705.1729 [hep-ph]. [Erratum: JHEP 11, 101 (2013)].

[164] **NA48/2** Collaboration, "Measurement of the form factors of charged kaon semileptonic decays," *JHEP* **10** (2018) 150, arXiv:1808.09041 [hep-ex].

[165] **GEANT4** Collaboration, S. Agostinelli *et al.*, "GEANT4–a simulation toolkit," *Nucl. Instrum. Meth. A* **506** (2003) 250–303.

[166] C. Gatti, "Monte Carlo simulation for radiative kaon decays," *Eur. Phys. J. C* **45** (2006) 417–420, arXiv:hep-ph/0507280.

[167] E. Goudzovski, "HNL production and exotics searches at NA62." https://indi.to/LYd7w. Talk at the Kaon 2019 conference.

[168] T. Yamazaki *et al.*, "Search for Heavy Neutrinos in Kaon Decay,".

[169] **NA62** Collaboration, "ADDENDUM I TO P326 Continuation of the physics programme of the NA62 experiment," Tech. Rep. CERN-SPSC-2019-039. SPSC-P-326-ADD-1. `https://cds.cern.ch/record/2691873`.

[170] **NA62** Collaboration, "2020 NA62 Status Report to the CERN SPSC," Tech. Rep. CERN-SPSC-2020-007. SPSC-SR-266. `http://cds.cern.ch/record/2713499`.

[171] P. Ballett, M. Hostert, and S. Pascoli, "Dark Neutrinos and a Three Portal Connection to the Standard Model," *Phys. Rev. D* **101** no. 11, (2020) 115025, `arXiv:1903.07589 [hep-ph]`.

[172] M. Drewes, "The Phenomenology of Right Handed Neutrinos," *Int. J. Mod. Phys.* **E22** (2013) 1330019, `arXiv:1303.6912 [hep-ph]`.

[173] F. F. Deppisch, P. S. Bhupal Dev, and A. Pilaftsis, "Neutrinos and Collider Physics," *New J. Phys.* **17** no. 7, (2015) 075019, `arXiv:1502.06541 [hep-ph]`.

[174] J. C. Pati and A. Salam, "Lepton Number as the Fourth Color," *Phys. Rev.* **D10** (1974) 275–289. [Erratum: Phys. Rev.D11,703(1975)].

[175] R. N. Mohapatra and J. C. Pati, "A Natural Left-Right Symmetry," *Phys. Rev.* **D11** (1975) 2558.

[176] R. N. Mohapatra and J. C. Pati, "Left-Right Gauge Symmetry and an Isoconjugate Model of CP Violation," *Phys. Rev.* **D11** (1975) 566–571.

[177] G. Senjanovic and R. N. Mohapatra, "Exact Left-Right Symmetry and Spontaneous Violation of Parity," *Phys. Rev.* **D12** (1975) 1502.

[178] M. Fukugita and T. Yanagida, "Resurrection of grand unified theory baryogenesis," *Phys. Rev. Lett.* **89** (2002) 131602, `arXiv:hep-ph/0203194 [hep-ph]`.

[179] S. Davidson, E. Nardi, and Y. Nir, "Leptogenesis," *Phys. Rept.* **466** (2008) 105–177, `arXiv:0802.2962 [hep-ph]`.

[180] M. Shaposhnikov, "Baryogenesis," *J. Phys. Conf. Ser.* **171** (2009) 012005.

[181] A. Pilaftsis, "The Little Review on Leptogenesis," *J. Phys. Conf. Ser.* **171** (2009) 012017, `arXiv:0904.1182 [hep-ph]`.

[182] L. Canetti and M. Shaposhnikov, "Baryon Asymmetry of the Universe in the νMSM," *JCAP* **1009** (2010) 001, `arXiv:1006.0133 [hep-ph]`.

[183] M. Drewes, B. Garbrecht, P. Hernandez, M. Kekic, J. Lopez-Pavon, J. Racker, N. Rius, J. Salvado, and D. Teresi, "ARS Leptogenesis," *Int. J. Mod. Phys.* **A33** no. 05n06, (2018) 1842002, `arXiv:1711.02862 [hep-ph]`.

[184] J. Klarić, M. Shaposhnikov, and I. Timiryasov, "Uniting low-scale leptogeneses," `arXiv:2008.13771 [hep-ph]`.

[185] J. Ghiglieri and M. Laine, "Sterile neutrino dark matter via coinciding resonances," *JCAP* **07** (2020) 012, arXiv:2004.10766 [hep-ph].

[186] A. Atre, T. Han, S. Pascoli, and B. Zhang, "The Search for Heavy Majorana Neutrinos," *JHEP* **05** (2009) 030, arXiv:0901.3589 [hep-ph].

[187] D. A. Bryman and R. Shrock, "Constraints on Sterile Neutrinos in the MeV to GeV Mass Range," *Phys. Rev. D* **100** (2019) 073011, arXiv:1909.11198 [hep-ph].

[188] **Belle** Collaboration, D. Liventsev *et al.*, "Search for heavy neutrinos at Belle," *Phys. Rev.* **D87** no. 7, (2013) 071102, arXiv:1301.1105 [hep-ex]. [Erratum: Phys. Rev.D95,no.9,099903(2017)].

[189] **E949** Collaboration, "Search for heavy neutrinos in $K^+ \to \mu^+ \nu_H$ decays," *Phys. Rev.* **D91** no. 5, (2015) 052001, arXiv:1411.3963 [hep-ex]. [Erratum: Phys. Rev.D91,no.5,059903(2015)].

[190] **NA62** Collaboration, "Search for heavy neutral lepton production in K^+ decays," *Phys. Lett.* **B778** (2018) 137–145, arXiv:1712.00297 [hep-ex].

[191] A. Izmaylov and S. Suvorov, "Search for heavy neutrinos in the ND280 near detector of the T2K experiment," *Phys. Part. Nucl.* **48** no. 6, (2017) 984–986.

[192] **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** no. 22, (2018) 221801, arXiv:1802.02965 [hep-ex].

[193] **NA62** Collaboration, "Search for K^+ decays to a muon and invisible particles," arXiv:2101.12304 [hep-ex].

[194] K. Cranmer and I. Yavin, "RECAST: Extending the Impact of Existing Analyses," *JHEP* **04** (2011) 038, arXiv:1010.2506 [hep-ex].

[195] A. Buckley, J. Butterworth, D. Grellscheid, H. Hoeth, L. Lonnblad, J. Monk, H. Schulz, and F. Siegert, "Rivet user manual," *Comput. Phys. Commun.* **184** (2013) 2803–2819, arXiv:1003.0694 [hep-ph].

[196] E. Conte, B. Fuks, and G. Serret, "MadAnalysis 5, A User-Friendly Framework for Collider Phenomenology," *Comput. Phys. Commun.* **184** (2013) 222–256, arXiv:1206.1599 [hep-ph].

[197] E. Conte and B. Fuks, "MadAnalysis 5: status and new developments," *J. Phys. Conf. Ser.* **523** (2014) 012032, arXiv:1309.7831 [hep-ph].

[198] M. Drees, H. Dreiner, D. Schmeier, J. Tattersall, and J. S. Kim, "CheckMATE: Confronting your Favourite New Physics Model with LHC Data," *Comput. Phys. Commun.* **187** (2015) 227–265, arXiv:1312.2591 [hep-ph].

[199] S. Kraml, S. Kulkarni, U. Laa, A. Lessa, W. Magerl, D. Proschofsky-Spindler, and W. Waltenberger, "SModelS: a tool for interpreting simplified-model results from the LHC and its application to supersymmetry," *Eur. Phys. J.* **C74** (2014) 2868, arXiv:1312.4175 [hep-ph].

[200] M. Papucci, K. Sakurai, A. Weiler, and L. Zeune, "Fastlim: a fast LHC limit calculator," *Eur. Phys. J.* **C74** no. 11, (2014) 3163, arXiv:1402.0492 [hep-ph].

[201] D. Barducci, A. Belyaev, M. Buchkremer, G. Cacciapaglia, A. Deandrea, S. De Curtis, J. Marrouche, S. Moretti, and L. Panizzi, "Framework for Model Independent Analyses of Multiple Extra Quark Scenarios," *JHEP* **12** (2014) 080, arXiv:1405.0737 [hep-ph].

[202] **LHC Reinterpretation Forum** Collaboration, W. Abdallah *et al.*, "Reinterpretation of LHC Results for New Physics: Status and Recommendations after Run 2," *SciPost Phys* **9** (3, 2020) 022, arXiv:2003.07868 [hep-ph].

[203] A. Abada, G. Arcadi, V. Domcke, M. Drewes, J. Klaric, and M. Lucente, "Low-scale leptogenesis with three heavy neutrinos," *JHEP* **01** (2019) 164, arXiv:1810.12463 [hep-ph].

[204] M. Chrzaszcz, M. Drewes, T. E. Gonzalo, J. Harz, S. Krishnamurthy, and C. Weniger, "A frequentist analysis of three right-handed neutrinos with GAMBIT," *Eur. Phys. J.* **C80** no. 6, (2020) 569, arXiv:1908.02302 [hep-ph].

[205] S. Davidson, G. Isidori, and A. Strumia, "The Smallest Neutrino Mass," *Phys. Lett. B* **646** (2007) 100–104, arXiv:hep-ph/0611389.

[206] I. Esteban, M. C. Gonzalez-Garcia, A. Hernandez-Cabezudo, M. Maltoni, and T. Schwetz, "Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23}, δ_{CP}, and the mass ordering," *JHEP* **01** (2019) 106, arXiv:1811.05487 [hep-ph].

[207] **APPEC Committee** Collaboration, A. Giuliani, J. J. Gomez Cadenas, S. Pascoli, E. Previtali, R. Saakyan, K. Schäffner, and S. Schönert, "Double Beta Decay APPEC Committee Report," arXiv:1910.04688 [hep-ex].

[208] **T2K** Collaboration, "Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations," *Nature* **580** no. 7803, (2020) 339–344, arXiv:1910.03887 [hep-ex]. [erratum: Nature583,no.7814,E16(2020)].

[209] **DUNE** Collaboration, "The DUNE Far Detector Interim Design Report Volume 1: Physics, Technology and Strategies," arXiv:1807.10334 [physics.ins-det].

[210] D. Wyler and L. Wolfenstein, "Massless Neutrinos in Left-Right Symmetric Models," *Nucl. Phys. B* **218** (1983) 205–214.

[211] G. Branco, W. Grimus, and L. Lavoura, "The Seesaw Mechanism in the Presence of a Conserved Lepton Number," *Nucl. Phys. B* **312** (1989) 492–508.

[212] M. Gonzalez-Garcia and J. Valle, "Fast Decaying Neutrinos and Observable Flavor Violation in a New Class of Majoron Models," *Phys. Lett. B* **216** (1989) 360–366.

[213] A. Abada, C. Biggio, F. Bonnet, M. B. Gavela, and T. Hambye, "Low energy effects of neutrino masses," *JHEP* **12** (2007) 061, arXiv:0707.4058 [hep-ph].

[214] A. Roy and M. Shaposhnikov, "Resonant production of the sterile neutrino dark matter and fine-tunings in the νMSM," *Phys. Rev. D* **82** (2010) 056014, arXiv:1006.4008 [hep-ph].

[215] M. B. Gavela, T. Hambye, D. Hernandez, and P. Hernandez, "Minimal Flavour Seesaw Models," *JHEP* **09** (2009) 038, arXiv:0906.1461 [hep-ph].

[216] 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," *JHEP* **07** (2014) 079, arXiv:1405.0301 [hep-ph].

[217] D. Alva, T. Han, and R. Ruiz, "Heavy Majorana neutrinos from $W\gamma$ fusion at hadron colliders," *JHEP* **02** (2015) 072, arXiv:1411.7305 [hep-ph].

[218] C. Degrande, O. Mattelaer, R. Ruiz, and J. Turner, "Fully-Automated Precision Predictions for Heavy Neutrino Production Mechanisms at Hadron Colliders," *Phys. Rev.* **D94** no. 5, (2016) 053002, arXiv:1602.06957 [hep-ph].

[219] **ATLAS** Collaboration, "Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data," *JINST* **14** no. 12, (2019) P12006, arXiv:1908.00005 [hep-ex].

[220] **ATLAS** Collaboration, "Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV," arXiv:2012.00578 [hep-ex].

[221] A. L. Read, "Presentation of search results: The CL_s technique," *J. Phys. G* **28** (2002) 2693–2704.