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

J.-L. Tastet,^{a,b} O. Ruchayskiy^a and I. Timiryasov^b

^aNiels Bohr Institute, University of Copenhagen,
Blegdamsvej 17, DK-2010, Copenhagen, Denmark

^bInstitute of Physics, Laboratory for Particle Physics and Cosmology,
École Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland

E-mail: jean-loup.tastet@epfl.ch, Oleg.Ruchayskiy@nbi.ku.dk,
Iar.Timiryasov@epfl.ch

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. However, the resulting exclusion limits may 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 a reinterpretation of the latest ATLAS search for HNLs decaying promptly to a tri-lepton final state. We show that in a realistic model with two HNLs, the actual limits can vary by several orders of magnitude depending on the free parameters of the model. Marginalizing over the 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. [1]. 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.

KEYWORDS: Beyond Standard Model, Neutrino Physics

ARXIV EPRINT: [2107.12980](https://arxiv.org/abs/2107.12980)

Contents

1	Introduction	1
1.1	Heavy neutral leptons	1
1.2	Motivation for a reinterpretation	2
2	Realistic neutrino oscillation models	4
2.1	The Lagrangian of the model	4
2.2	Parametrization of the Yukawas	4
2.3	Heavy neutrino mixing	5
2.4	Quasi-Dirac HNLs, lepton number violating effects and relevant limits	7
3	Procedure	8
3.1	Event selection	8
3.2	Signal	10
3.2.1	MadGraph setup	10
3.2.2	Signal computation for arbitrary model parameters	11
3.2.3	Signal computation for quasi-Dirac HNLs	12
3.3	Efficiencies	13
3.4	Background	17
3.5	Statistical limits	18
3.5.1	CL_s technique: a general reminder	18
3.5.2	CL_s technique: implementation	19
4	Results	20
4.1	Majorana-like HNL pair	22
4.2	Dirac-like HNL pair	25
5	Conclusion & outlook	26
5.1	Reinterpretation	26
5.2	Wish-list for a painless reinterpretation of future experimental results	28
A	Ancillary files	30
A.1	Card files for the Monte-Carlo event generation	30
A.2	Signal cross sections	30
A.3	Signal efficiencies	30

1 Introduction

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. 2–5]. 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. [2, 3, 6–10] 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_μ^2, U_τ^2). HNLs first appeared in the context of left-right symmetric models [11–14] which required an extension of the fermion sector with Standard Model (SM) gauge singlet particles, and then in the (type I) see-saw mechanism [15–22] 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 [23]. This scenario (known as *leptogenesis*) has been actively developed since the 1980s (see reviews [24, 25]). 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 [7, 26–30]; for a recent overview see e.g. [31, 32]. 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 [6–8, 33, 34] within the *Neutrino Minimal Standard Model* (ν MSM).

Starting from the 1980s [35–38], many experiments have searched for HNLs (as summarized e.g. in refs. [3, 4, 10, 39–41]). Current generation particle physics experiments, including LHCb, CMS, ATLAS, T2K, Belle and NA62, all include HNL searches into their scientific programs [1, 42–53]. However, as pointed out in ref. [54], 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 it 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. 4], with the notable exception of the SHiP experiment, which provided sensitivity estimates for arbitrary sets of mixing angles [55]. This raises a few questions:

1. Which 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. [56–63], that allow recasting LHC results for new sets of models (see also [64]). These tools have mostly been developed in the context of supersymmetry and similar searches at the LHC and are not readily applicable to HNL models, whose collider phenomenology is quite different.

Figure 1. Lepton number conserving (LNC) and violating (LNV) diagrams contributing to the same $\mu^+\mu^+e^-$ + missing transverse energy (MET) final state.

In this work we perform a step in the direction of *recasting LHC results*. Specifically, we recast the ATLAS tri-lepton search [1] 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.¹ In what follows we will refer to this model as a *realistic HNL model*. As we shall see below, even in this simple model, the interpretation of the results is a non-trivial task.

1.2 Motivation for a reinterpretation

The realistic seesaw model describing neutrino oscillations brings several changes compared to the single-HNL, single-flavor model analyzed by the ATLAS collaboration [1]. The analysis from ref. [1] concentrated on the following process:

$$pp \rightarrow W^\pm + X \quad \text{with} \quad W^\pm \rightarrow \ell_\alpha^\pm + N \quad \text{followed by} \quad N \rightarrow \ell_\alpha^\pm + \ell_\beta^\mp + \bar{\nu}_\beta^{(-)} \quad (1.1)$$

where $\ell_\alpha^\pm$ are light leptons ($e^\pm$ or $\mu^\pm$), $\alpha \neq \beta$ and $\bar{\nu}_\beta^{(-)}$ is a neutrino or anti-neutrino with flavor β . They 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 1), along with its CP-conjugate, contributed to the final signal. The upper limit on an admissible signal was thus 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:

¹In the case of three or more HNLs, the constraints on the HNL mixing angles are much more relaxed thanks to the freedom conferred by the additional model parameters [65, 66]. Such models can thus accommodate more extreme ratios of the mixing angles (and, for four or more HNLs, even allow some mixing angles to be zero). However, many of the results that we will discuss in this paper still apply. In particular, most points in the parameter space of these models also correspond to non-trivial mixing patterns, and in order to probe them the ATLAS results will need to be recast.

1. In such a model, several processes contribute incoherently² 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_μ^2, U_τ^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 take into account that the signal is non-zero in both channels.
3. Different processes that contribute to the same tri-lepton final state have different kinematics (due in part to spin correlations [67]). 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. [67] or [68–74] 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. [70, 75–77]. Since different processes (such as those in figure 1) have different kinematics and thus efficiencies, this implies that 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” limits (which we will define in section 2).

All these points make it impossible to reinterpret the ATLAS results by just rescaling them (as done e.g. in ref. [78]). Instead one should perform a full signal and background modeling and evaluate the signal selection efficiencies. Although this can only be done properly by the collaboration itself, thanks to 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 recommendations from the LHC Reinterpretation Forum [64].

The present paper is organized as follows. In section 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 constrain its parameters, and examine how interference effects between multiple HNLs can completely change their phenomenology. We then describe our analysis procedure in section 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, 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 5, and summarize what data should ideally be reported by experiments in order to allow reinterpreting their limits easily and accurately within realistic models.

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

2 Realistic neutrino oscillation models

2.1 The Lagrangian of the model

Our starting point is the type I seesaw mechanism [15, 16, 18, 19, 21, 22], that we briefly review below. The exposition is fairly standard and can be found, e.g. in refs. [3, 8, 33] and [79, ch. 14]. The reader can skip it, taking notice of the definitions (2.2)–(2.3).

The Lagrangian of the model reads

$$\mathcal{L}_{\text{SM+HNL}} = \mathcal{L}_{\text{SM}} + i\bar{\nu}_{R_I}\not{\partial}\nu_{R_I} - F_{\alpha I} \left(\bar{L}_\alpha \cdot \tilde{\Phi} \right) \nu_{R_I} - \frac{1}{2} M_I \bar{\nu}_{R_I}^c \nu_{R_I}, \quad (2.1)$$

where $\mathcal{L}_{\text{SM}}$ is the usual SM Lagrangian and 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 (2.1) obtains a vacuum expectation value $\langle \Phi \rangle = (0, v)^T$ and the Yukawa interaction terms in eq. (2.1) 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

$$\nu_{L\alpha} = V_{\alpha i}^{\text{PMNS}} \nu_i + \Theta_{\alpha I} N_I^c, \quad (2.2)$$

where V^{PMNS} is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix (see e.g. [81]). 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

$$\Theta_{\alpha I} \simeq \frac{v F_{\alpha I}}{M_I}. \quad (2.3)$$

2.2 Parametrization of the Yukawas

The Lagrangian (2.1) contains 11 new parameters, as compared to the SM one [3]. These parameters are, however, constrained by neutrino oscillation data [82]. Five neutrino parameters have already been measured: two mass differences (Δm_{atm}^2 and Δm_{sun}^2) 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

³Given the Lagrangian (2.1) with two right-handed neutrinos, the lightest neutrino is massless (up to quantum corrections [80]).

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 [86]:

$$F = \frac{i}{v} V^{\text{PMNS}} \sqrt{m_\nu^{\text{diag}}} R \sqrt{M^{\text{diag}}}, \quad (2.4)$$

where the matrix $M^{\text{diag}} = \text{diag}(M_1, M_2)$, and R is a complex 3×2 matrix satisfying $R^T R = 1_{2 \times 2}$. For the PMNS matrix we use the standard parametrization [79]. 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 $m_3 = 0$ for IH.

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

$$R^{\text{NH}} = \begin{pmatrix} 0 & 0 \\ \cos \omega & \sin \omega \\ -\xi \sin \omega & \xi \cos \omega \end{pmatrix}, \quad R^{\text{IH}} = \begin{pmatrix} \cos \omega & \sin \omega \\ -\xi \sin \omega & \xi \cos \omega \\ 0 & 0 \end{pmatrix}. \quad (2.5)$$

with a complex angle $\omega = \text{Re } \omega + i \text{Im } \omega$, and a discrete parameter $\xi = \pm 1$. Changing the sign of ξ can be undone by $\omega \rightarrow -\omega$ along with $N_2 \rightarrow -N_1$ [87], so we fix $\xi = +1$.

2.3 Heavy neutrino mixing

The weak-like interactions of HNLs are suppressed by the *mixing angles* $\Theta_{\alpha I}$ defined in eq. (2.3). 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

$$U_\alpha^2 \equiv \sum_I |\Theta_{\alpha I}|^2 \quad \text{and} \quad U_{\text{tot}}^2 \equiv \sum_{\alpha, I} |\Theta_{\alpha I}|^2, \quad (2.6)$$

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:

$$U_{\text{tot}}^2 = \frac{\sum_i m_i}{M_N} \cosh(2 \text{Im } \omega) \quad (2.7)$$

where $M_N = \frac{1}{2}(M_1 + M_2)$. For the corresponding expressions of U_α^2 , see e.g. ref. [88].

⁴These parameters may be probed in a not-so-distant future: for the inverted hierarchy, the next generation of neutrinoless double beta decay experiments may provide information on the Majorana phases [83], while the CP -violating phase δ is already constrained by T2K [84], with further improvements expected from the DUNE experiment [85].

⁵These complex phases can be important if *the period of HNL oscillations* is comparable with the size of the experiment, see e.g. [67] and references therein.

Figure 2. Ternary plot showing the combinations of mixing angles $U_\alpha^2/U_{\text{tot}}^2$, $\alpha = e, \mu, \tau$, which are consistent with the NuFIT 5.0 [92, 93] 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.

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 $|\text{Im } \omega| \gg 1$ 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

$$U_e^2/U_{\text{tot}}^2 + U_\mu^2/U_{\text{tot}}^2 + U_\tau^2/U_{\text{tot}}^2 = 1, \quad (2.8)$$

we can display the combinations of U_α^2 which are compatible with neutrino oscillation data using a ternary plot as in figure 2, cf. [78, 89–91]. In our analysis, we used the most recent global fit to neutrino oscillation data, NuFIT 5.0 [92, 93]. 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 2.⁶ In order to better visualize the

⁶We have used the $\Delta\chi^2$ tables from NuFIT 5.0 which do not include the atmospheric data provided by the Super-Kamiokande collaboration [94]. Our choice of benchmark models is only slightly affected by this choice, and this does not qualitatively change our analysis or conclusions.

correspondence between the exclusion limits and various points in the allowed regions, we have defined a number of benchmarks, which are represented in figure 2.

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.

$$\Delta M \equiv |M_1 - M_2| \ll M_N = \frac{M_1 + M_2}{2}. \quad (2.9)$$

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 [75–77, 95–101]. Secondly, low-scale leptogenesis (see the recent work [31] 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 [28, 34, 102].

In the limit $M_1 \approx M_2$ there is an approximate global U(1) symmetry in the theory.⁷ In this *quasi-Dirac limit* of the two-HNLs model, the lepton number violating (LNV) processes (such as 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. [7, 67–74]. 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 $\tau = \sqrt{(x_D - x_P)^2}$ between the HNL production and decay vertices [67]:

$$d\Gamma_{\alpha\beta}^{\text{LNC/LNV}}(\tau) \cong 2 |\Theta_{\alpha 1}|^2 |\Theta_{\beta 1}|^2 \left(1 \pm \cos(\Delta M \tau)\right) e^{-\Gamma \tau} d\hat{\Gamma}_{\alpha\beta}^{\text{LNC/LNV}} \quad (2.10)$$

with the (+) sign for LNC and (−) for LNV, and where $d\hat{\Gamma}_{\alpha\beta}^{\text{LNV/LNC}}$ 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 $\Gamma \stackrel{\text{def}}{=} \Gamma_1 \cong \Gamma_2$ 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 (2.10) demonstrates the two limiting cases of the two-HNLs seesaw model:

$$\begin{aligned} \Delta M \tau \ll 2\pi \quad (\text{Dirac-like limit}) \quad & d\Gamma_{\alpha\beta}^{\text{LNV}} \approx 0, \quad d\Gamma_{\alpha\beta}^{\text{LNC}} \text{ is enhanced by } \cong 2 \\ \Delta M \tau \gg 2\pi \quad (\text{Majorana-like limit}) \quad & \text{integrated partial widths } \Gamma_{\alpha\beta}^{\text{LNV}} \cong \Gamma_{\alpha\beta}^{\text{LNC}} \end{aligned} \quad (2.11)$$

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

⁷The 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 $\frac{1+\gamma_5}{2} \Psi = \frac{\nu_{R1} + i \nu_{R2}}{\sqrt{2}}$.

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 .

3 Procedure

In order to reinterpret the limits from the ATLAS prompt search [1] (with extra details in the Ph.D. thesis [103]) we have tried to reproduce the ATLAS analysis as accurately as possible. Our signal is simulated using MADGRAPH5_AMC@NLO [104] with the HEAVYN model [105, 106] (section 3.2). For the event selection (section 3.1), we have implemented the ATLAS cut flow and obtained comparable efficiencies (section 3.3). We take the total background counts from the ATLAS publication [1] (section 3.4). Finally, in order to compute the limits (section 3.5), we use the CL_s test statistics, along with a very simplified treatment of uncertainties.

3.1 Event selection

The prompt ATLAS analysis [1] 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 analysis⁸ 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 cut flow allows reproducing the ATLAS efficiencies with sufficient accuracy for the purpose of this reinterpretation. The list of cuts is shown in table 1, and their order roughly follows that of ref. [103]. When different cuts were applied to the 2015 and 2016 datasets, we use the 2016 cuts, since the 2015 dataset is smaller than the 2016 one by about an order of magnitude.

1. We start by applying a cut on the distance of closest approach to the origin in the r - z plane: $|\Delta z_0 \sin(\theta)| < 0.5$ mm for the leading lepton⁹ and $|\Delta z_0 \sin(\theta)| < 1$ mm for the remaining ones.

⁸By “process” we mean a set of diagrams which have the same incoming and outgoing particles (distinguishing in particular all possible flavors of outgoing (anti-)neutrinos). Thus, e.g., figure 1 corresponds to two different *processes* in our terminology, although they share the same visible final state. The same notion is used in MadGraph.

⁹In the electron channel, the leading lepton is defined as the electron with the highest p_T , and in the muon channel as the muon with the highest p_T .

2. Next, we apply the default transverse momentum and pseudorapidity requirements on the three changed leptons, i.e. $p_T > 4.5 \text{ GeV}$ and $|\eta| \in [0, 1.37 \cup 1.52, 2.47[$ for all electrons and $p_T > 4 \text{ GeV}$ and $|\eta| < 2.5$ for all muons.
3. Then, we simulate the selection performed by the trigger by applying the relevant p_T requirements, as found in ref. [1], ch. 4.1, §1. For the single-electron trigger used in the electron channel, we do not apply the ID requirements, since the ID efficiency is difficult to accurately estimate.
4. We then apply the trigger offline requirements on the two leading leptons: $p_T(e_{\text{lead}}) > 27 \text{ GeV}$ and $p_T(e_{\text{sublead}}) > 10 \text{ GeV}$ for the electron channel and $p_T(\mu_{\text{lead}}) > 23 \text{ GeV}$ and $p_T(\mu_{\text{sublead}}) > 14 \text{ GeV}$ for the muon channel.
5. Next, we require the tri-lepton invariant mass M_{3l} to be in the interval $]40, 90[\text{ GeV}$.
6. We then 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. [107] for electrons and ref. [108] 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, $M(e, e) < 78 \text{ GeV}$, in order to veto the background from $Z \rightarrow e^+ e^-$ where one of the electron charges is misreconstructed.
8. The missing transverse energy is then restricted to $E_T^{\text{miss}} < 60 \text{ GeV}$.

Finally, the events passing the above cuts are binned in $M(l_{\text{sublead}}, l')$, 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 $[40, 50[\text{ GeV}$.

Our cut flow is summarized in table 1. One notable difference with the ATLAS paper is the absence of a b -jet veto in our analysis, which we omitted since b -jets appear in only $\mathcal{O}(1\%)$ of the signal events, therefore this cut would remove almost no signal *at truth level*. For this reason, we have not generated b -jets in our final samples. A further difference comes from the cuts related to the displacement of the leading lepton. ATLAS additionally imposes $|d_0/\sigma(d_0)| < 5$ (electron) or < 3 (muon), while we only impose the $|\Delta z_0 \sin(\theta)|$ cut and omit the d_0 cut, since we do not know $\sigma(d_0)$ well enough.¹⁰ This most likely does not affect the signal at truth level, since the leading lepton has a very small displacement in all relevant cases: 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. We also decided to omit the lepton identification (ID) requirements, whose efficiency is harder to model for electrons due to being significantly less smooth [107] than the isolation one, in particular for the “tight” working point and at low p_T . For muons the ID efficiency is close enough to 1 [108] that it can probably be safely neglected. Our attempt at implementing this cut only resulted in a significantly decreased accuracy for the efficiency estimates. A possible cause could be that the tabulated efficiencies have been computed using different sets of triggers and cuts and therefore cannot be transposed directly to the present analysis.

¹⁰In principle, the data on $\sigma(d_0)$ is reported in ref. [109] (figure 4). However, it only exists for muons and is too coarse-grained to be exploitable in our analysis.

#	Electron channel	Muon channel
1	$ \Delta z_0 \sin(\theta) (l) < \begin{cases} 0.5 \text{ mm for } l = l_{\text{lead}} \\ 1 \text{ mm for } l \neq l_{\text{lead}} \end{cases}$	
2	$p_T(e) > 4.5 \text{ GeV}^\dagger, p_T(\mu) > 4 \text{ GeV}$ $ \eta(e) \in [0, 1.37] \cup [1.52, 2.47[, \eta(\mu) < 2.5$	
3	$p_T(e_{\text{lead}}) > 26 \text{ GeV}^\dagger$	$p_T(\mu_{\text{lead}}) > 22 \text{ GeV}^\dagger$ AND $p_T(\mu_{\text{sublead}}) > 8 \text{ GeV}$
4	$p_T(e_{\text{lead}}) > 27 \text{ GeV}$ $p_T(e_{\text{sublead}}) > 10 \text{ GeV}$	$p_T(\mu_{\text{lead}}) > 23 \text{ GeV}$ $p_T(\mu_{\text{sublead}}) > 14 \text{ GeV}$
5	$40 \text{ GeV} < M(l, l') < 90 \text{ GeV}$	
6	“Loose” lepton isolation	
7	$Z \text{ veto: } M(e, e) < 78 \text{ GeV}$	—
8	$E_T^{\text{miss}} < 60 \text{ GeV}$	

Table 1. Our cut flow for the electron and muon channels. The † indicates cuts which differ between 2015 and 2016 (the 2016 cuts were used in this analysis). Lepton identification and $|d_0/\sigma(d_0)|$ cuts have been omitted due to the complexity of their implementation.

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 $M(l_{\text{sublead}}, l')$ bin in each signal region, for any model parameters. We do so using a simple model, described below.

3.2.1 MadGraph setup

The signal processes contributing to each channel are listed in tables 2 and 3.¹¹ 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).

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 ~ 40000 weighted events generated at leading order using MADGRAPH5_AMC@NLO v2.8.x [104] along with the HEAVYN model [105, 106] (specifically, we use the SM_HEAVYN_CKM_ALLMASSES_LO model¹²), which includes the non-diagonal CKM

¹¹In 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.

¹²Note that HNLs are Majorana particles in this model. This is actually not a problem for simulating quasi-Dirac HNLs: all we need to do is suitably rescale or suppress the cross section of each process, as will be discussed in section 3.2.3.

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

Table 2. Signal processes contributing to the electron channel. Up to two additional hard jets have been included in the process string, but are omitted here for brevity.

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

Table 3. Signal processes contributing to the muon channel. Up to two additional hard jets have been included in the process string, but are omitted here for brevity.

matrix as well as the finite fermion masses. The center of mass energy is set to $\sqrt{s} = 13$ 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 2 and 3, with up to two additional hard jets (excluding b -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.

3.2.2 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 angles¹³ $|\Theta_e|$, $|\Theta_\mu|$ 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 it for each new set of mixing angles. This is done as follows.

As a first step, we generate Monte-Carlo samples for all the processes listed in tables 2 and 3, for each HNL mass $M_N \in \{5, 10, 20, 30, 50\}$ GeV and using the reference parameters $|\Theta|_{\text{ref}} = 10^{-3}$ and $\Gamma_{\text{ref}} = 10^{-5}$ GeV as placeholders for the remaining model parameters.¹⁴ For each process P , we only set the relevant mixing angle $|\Theta_{\alpha(P)}|$ and $|\Theta_{\beta(P)}|$ to $|\Theta|_{\text{ref}}$,

¹³Since we are dealing with 2 HNLs far from the seesaw line, $\Theta_{\alpha 2} \cong \pm i \Theta_{\alpha 1}$ [75, 76]. 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. (2.6).

¹⁴These parameters allow for the successful numerical integration in the narrow width approximation.

where $\alpha(P)$ and $\beta(P)$ respectively correspond to the generations coupling to the HNL at production and decay, as listed in tables 2 and 3.

The key observation here is that the branching fraction of $W^+ \rightarrow l_\alpha N$ is proportional to $|\Theta_\alpha|^2$, while the branching fraction of $N \rightarrow l_\beta l_\gamma \nu_\gamma$ 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:

$$\sigma_P(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) = \sigma_P^{\text{ref}} \times \frac{|\Theta_{\alpha(P)}|^2 |\Theta_{\beta(P)}|^2}{|\Theta|_{\text{ref}}^4} \times \frac{\Gamma_{\text{ref}}}{\Gamma_N(M_N, \Theta_e, \Theta_\mu, \Theta_\tau)} \quad (3.1)$$

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:

$$\tau_N^{-1} = \Gamma_N(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) = \sum_{\beta=e,\mu,\tau} |\Theta_\beta|^2 \times \hat{\Gamma}_\beta(M_N) \quad (3.2)$$

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: $N_P = \mathcal{L}_{\text{int}} \times \sigma_P$.

3.2.3 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 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.

Nature	$c_P, P \in \text{LNC}$	$c_P, P \in \text{LNV}$	$c_\Gamma = \Gamma_N/\Gamma_{\text{Maj.}}$
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 ¹⁵	4	0	1

Table 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.

3.3 Efficiencies

In order to obtain a sensitivity estimate, we must compute the expected signal count in every $M(l_{\text{sublead}}, l)$ bin reported by the ATLAS collaboration.¹⁶ This 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 of the HNL. 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:

$$N_b = \mathcal{L}_{\text{int}} \times \sum_P \epsilon_{P,b}(M_N, \tau_N) \times c_P \times \sigma_P(M_N, \Theta_e, \Theta_\mu, \Theta_\tau) \quad (3.3)$$

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 3.1 and table 1. The binned efficiency is then:

$$\epsilon_{P,b} = \frac{\sum (\text{weights of events after cuts, which end up in bin } b)}{\sum (\text{weights of all events before cuts, from any bin})} \quad (3.4)$$

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.¹⁷ Similarly, we can obtain the unbinned efficiency as:

$$\epsilon_P = \frac{\sum (\text{all event weights after cuts})}{\sum (\text{all event weights before cuts})}. \quad (3.5)$$

¹⁵Note how a quasi-Dirac pair in the $\delta M = 0$ limit is equivalent to a single Dirac HNL, up to a redefinition of the mixing angles $\Theta_{\alpha I}^{\text{quasi-Dirac}} = \Theta_{\alpha I}^{\text{Dirac}}/\sqrt{2}$.

¹⁶We 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.

¹⁷Some cuts (such as lepton ID and isolation cuts) are implemented by reweighting events using tabulated efficiencies.

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 μ m

Table 5. Benchmark points (taken from ref. [1]) used to plot the efficiencies in figures 3 and 4. Note that our calculations are more general, and work for any combination of M_N and τ_N .

The unbinned efficiencies for the four LNV processes are plotted in figure 3 along with the efficiencies reported by ATLAS in ref. [1], while those for LNC processes are plotted in figure 4. Since the efficiency of a process depends on both the HNL mass and its lifetime, we had to choose a set of benchmark points to produce figures 3 and 4. In order to be able to compare our efficiency calculation with the ATLAS efficiencies, we have chosen the same benchmarks as reported in ref. [1] and reproduced in table 5. Our estimate is reasonably accurate for the muon channel, with a mean relative error¹⁸ of 18% (maximum 48%), but less so for the electron channel, with a mean relative error of 38% and a factor of ~ 4 in the worst case (which corresponds to the lowest HNL mass hypothesis $M_N = 5$ GeV). The main difference between the two channels comes from the larger reliance on the electron ID (which we ignore) in the electron channel. Indeed, the electron ID is used for the single-electron trigger as well as for the ID cuts on both electrons; and contrary to the “loose” muon ID [108, figure 12] used for muons, its efficiency can be significantly smaller than 1 [107, figure 17]. This omission could contribute to the worse agreement between signal efficiencies in the electron channel. Another potential factor could be the large HNL displacement. The displacement has not been taken into account when tabulating the isolation efficiencies (computed for $Z \rightarrow \ell\ell$ in refs. [107, 108]). This would explain why the discrepancy is stronger for larger $c\tau_N\gamma_N$. Comparing figure 4 with figure 3, we also notice that **the efficiencies for LNC processes can be significantly smaller than for LNV processes**. This is mostly due to the different spin correlation patterns [67, 110] for LNC vs. LNV leading to different lepton spectra and to different geometrical acceptances of the lepton p_T and displacement cuts.

Since the original Monte-Carlo samples used for this analysis did not take spin correlations into account, and were generated under the single-flavor mixing hypothesis, the cut flow has been optimized under these assumptions. In principle, this might lead to a sub-optimal cut selection when it is applied to more realistic models (which we eventually hope to *observe*). For this reason, we would generally recommend performing the cut optimization using a set of signal samples which are representative of realistic models (instead of simplified benchmarks) and which have been generated using a Monte-Carlo event generator (such as MADGRAPH) which can model spin correlations. However, in the present case, it seems that most cuts were chosen solely based on the minimal requirements imposed by the existing triggers, and therefore would not have been very different had the cut optimization been performed with more realistic models in mind.

Even using the extrapolation method described above and eq. (3.3), one efficiency $\epsilon_{P,b}(M_N, \tau_N)$ must in principle still be computed for every process P , bin b , HNL mass

¹⁸We define the relative error on the total efficiency ϵ as $\frac{|\epsilon_{\text{ours}} - \epsilon_{\text{ATLAS}}|}{\frac{1}{2}(\epsilon_{\text{ours}} + \epsilon_{\text{ATLAS}})}$.

Figure 3. Cumulative unbinned signal efficiencies (for the total event count, i.e. summed over all bins) after applying each cut listed in table 1, computed for the benchmark points found in ref. [1]. The black dashed line denotes the total efficiencies reported in ref. [1], table 2, 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 only relevant processes in the original prompt search.

M_N and lifetime τ_N . However, several simplifications exist. First, the efficiencies for the full set of $M(l_{\text{sublead}}, l')$ 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:

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

The “simplest” functional form satisfying these two conditions is:

$$\epsilon(\tau_N) = \frac{\epsilon_0}{1 + \frac{\tau_N}{\tau_0}} \quad (3.6)$$

Figure 4. Cumulative unbinned signal efficiencies (for the total event count, i.e. summed over all bins) after applying each cut listed in table 1, computed for the benchmark points found in ref. [1], for lepton number *conserving* (LNC) processes. The gray line with diamond markers corresponds to the total efficiency.

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, 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 5 for both the binned and unbinned efficiencies, for the $W^+ \rightarrow e^+(N \rightarrow e^+ \mu^- \bar{\nu}_\mu)$ 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 (we used 13) 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

Figure 5. Binned and unbinned efficiencies as a function of the HNL lifetime τ_N , for the process $W^+ \rightarrow e^+(N \rightarrow e^+\mu^-\bar{\nu}_\mu)$ with $M_N = 30$ 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.

HNL pairs.¹⁹ Lifetime reweighting can additionally be used to simulate intermediate lifetimes without having to generate new samples. 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.

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 [1], 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.

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

¹⁹Plus three samples for computing the HNL lifetime, but these only need to be run at parton level and therefore have a negligible computational cost.

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. [1]. 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. [1]. Since the uncertainties on the individual components of the background are unfortunately not reported, implementing a statistical test necessarily requires some guessing on our side. After experimenting with several well-motivated background models and selecting the one which leads to the best approximation of the ATLAS limits, we have decided to model the uncertainty as being entirely caused by a single, Gaussian-constrained normalization factor μ_{tot} . In other words, we assume that the background expectations in the various $M(l_{\text{sublead}}, l')$ bins are maximally correlated. This is consistent with the observation that the statistical errors on the normalization factors μ_{mf} and $\mu_{t\bar{t}}$ are among the leading uncertainties. The accuracy of this simplified model will be explicitly tested in section 3.5.

3.5 Statistical limits

Ref. [1] 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 [111]. For completeness, a quick reminder about the CL_s technique follows in section 3.5.1. Knowledgeable users are welcome to skip it and go directly to section 3.5.2.

3.5.1 CL_s technique: a general reminder

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

$$t(x) \equiv 2 \ln \left(\frac{\mathcal{L}(x|H_{s+b})}{\mathcal{L}(x|H_b)} \right) \quad (3.7)$$

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:

$$\text{CL}_b = \mathcal{P}(t(X) < t_{\text{obs}} | H_b) = \int_{-\infty}^{t_{\text{obs}}} dt p_b(t) \quad (3.8)$$

$$\text{CL}_{s+b} = \mathcal{P}(t(X) < t_{\text{obs}} | H_{s+b}) = \int_{-\infty}^{t_{\text{obs}}} dt p_{s+b}(t) \quad (3.9)$$

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:

$$\text{CL}_s = \frac{\text{CL}_{s+b}}{\text{CL}_b} \in [0, 1] \quad (3.10)$$

and a given signal + background hypothesis H_{s+b} is considered to be excluded if $\text{CL}_s < 0.05$. For any signal stronger than the $\text{CL}_s = 0.05$ 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.

3.5.2 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 $M(l_{\text{sublead}}, l')$ 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 $i = 1 \dots 5$ for the electron channel and $i = 6 \dots 10$ for the muon channel. As discussed in section 3.4, we model the background as a set of expectation values $\{b_i\}$ for each bin $i = 1 \dots 10$ (taken from the ATLAS paper) along with a Gaussian-constrained normalization factor μ_{tot} with standard deviation $\sigma_{\text{tot}} = \sum_i (b_i^+ - b_i^-) / (2 \sum_i b_i)$, where the $-$ and $+$ superscripts respectively denote the lower and upper uncertainty bands from the ATLAS plot (see table 6). The signal is modeled as a set of signal expectations $\{s_i\}$, $i = 1 \dots 10$, which we compute for each set of the model parameters $(M_N, \Theta_e, \Theta_\mu, \Theta_\tau)$ using the method described in sections 3.2 and 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.²⁰ 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 $\mu_{\text{tot}} b_i$ for the background-only hypothesis and $\mu_{\text{tot}} b_i + s_i$ for the signal + background hypothesis. The full likelihood for the signal + background hypothesis is thus:

$$\begin{aligned} \mathcal{L}(x | H_{s+b}) &= \mathcal{P}(\mu_{\text{tot}} | \mathcal{N}(1, \sigma_{\text{tot}})) \times \prod_{i=1}^{10} \mathcal{P}(x_i | \text{Pois}(\mu_{\text{tot}} b_i + s_i)) \\ \text{where } \mu_{\text{tot}} &= \frac{\sum_i (x_i - s_i)}{\sum_i b_i} \end{aligned} \quad (3.11)$$

²⁰In 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.

$ee\mu$				$\mu\mu e$			
Expected background			Observed count	Expected background			Observed
b_i	b_i^-	b_i^+	x_i	b_i	b_i^-	b_i^+	x_i
19.0	14.9	23.1	19	21.3	17.5	25.1	23
18.0	14.4	21.7	20	13.8	10.6	17.0	15
21.0	17.4	24.7	19	18.7	15.1	22.3	20
13.6	10.9	16.2	15	13.3	10.6	16.1	14
6.1	4.2	7.8	5	13.1	10.2	15.9	13

Table 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 values have been digitized from figure 5 in [1]. Only the total background expectation (without the individual contributions) is shown.

The likelihood for the background-only hypothesis H_b is obtained by setting the signal s_i to zero in eq. (3.11).

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. [1]). 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 6. The mean ratio between our limits and the ones from ATLAS is 0.64, and the worst-case ratio is 0.42. 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 Results

Below we present our results — the exclusion limits for the model with two HNLs. We calculate exclusions for each of the benchmark points defined in figure 2. Benchmarks are chosen in such a way as to represent both typical and extreme ratios of the mixing angles $U_e^2 : U_\mu^2 : U_\tau^2$. As each benchmark fixes the mixing pattern, our results are most compactly expressed as exclusion limits for the total mixing angle $U_{\text{tot}}^2 = U_e^2 + U_\mu^2 + U_\tau^2$ (eq. (2.7)). Figures 7 and 8 present our results for the Majorana- and Dirac-like cases respectively. The limits for the flavor mixing angles U_α^2 are presented in figures 9 to 12. All these limits are the observed exclusion limits, and all of them (including the single-flavor limits) have been derived using the same statistical method,²¹ which we described in 3.5.

²¹In this way all the limits in figure 7 to figure 12 are obtained by means of the same statistical procedure and under the same assumptions about the systematic uncertainties. Therefore, although they might slightly deviate from the actual ATLAS limits from ref. [1], they should be comparable among themselves. For comparison of our limits with those derived by ATLAS, see figure 6.

Figure 6. 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 7. Original (black lines) and reinterpreted (colored lines) 95% exclusion limits on the total mixing angle $U_{\text{tot}}^2 = \sum_{\alpha=e,\mu,\tau} \sum_{I=1,2} |\Theta_{\alpha I}|^2$ for a *Majorana-like* HNL pair for the normal (left) and inverted (right) mass orderings. The black lines are limits obtained under the single-flavor assumption, while the solid colored lines denote those obtained for the benchmark points defined in figure 2. When scanning over all ratios of mixing angles allowed by neutrino oscillation data, the exclusion limits span the blue (green) shaded regions. Correspondingly, the gray filled area is excluded at $\text{CL} > 95\%$ for all possible ratios of mixing angles, and thus constitutes an exclusion limit independent of the specific choice of mixing angles, valid as long as we consider the two HNL model explaining neutrino oscillations.

The legend for these plots is as follows. The thick dashed and dotted lines in each plot represent the exclusion limits obtained under the assumption of a Majorana-like HNL pair mixing with a single flavor (respectively the electron and muon flavor). Up to a factor of 2, this corresponds to the scenario considered by ATLAS in the current prompt search [1]. These limits are grayed out in the plots for the Dirac-like pair in order to emphasize that the search has *no sensitivity* to the Dirac-like case for the single-flavor mixing. The solid colored lines denote the exclusion limits obtained for the various benchmark points defined in figure 2. The benchmarks can be identified using the numbers in the right margin. The colored, filled area represents the set of possible (“benchmark-dependent”)²² limits spanned by all the combinations of mixing angles allowed by the NuFIT 5.0 neutrino data (at 95% CL).²³ In other words, it shows the dependence of the exclusion limits on the specific combination of mixing angles, within the constraints from neutrino oscillation data (which are represented by the similarly-colored area in figure 2). Finally, the gray filled area denotes the set of mixing angles which are excluded at the 95% level for all the allowed ratios of mixing angles. It thus represents the most conservative (benchmark-independent) limit that can be obtained for a given model. No choice of mixing angles that is in agreement with neutrino oscillation data (within the 2 HNL seesaw model²⁴) can produce a limit within the gray filled region.

4.1 Majorana-like HNL pair

Let us first consider the case of a Majorana-like HNL pair, which is closer to the “single Majorana HNL” model considered by ATLAS and many other experiments. The relevant limits are shown in figures 7, 9 and 10. Apart from a trivial factor of two due to the two nearly degenerate mass eigenstates, the main difference with ATLAS is that in a realistic seesaw model the HNLs must mix with all three flavors at the same time. Looking at the total mixing angle in figure 7, we immediately notice that the limits on U_{tot}^2 are weaker than the single-flavor mixing limits for all our benchmarks, sometimes by more than an order of magnitude. The pattern is obvious for the normal hierarchy (but also visible for the inverted one): the benchmark points which have the strongest tau fraction $x_\tau = U_\tau^2/U_{\text{tot}}^2$ also have the worst sensitivity. This was already observed in ref. [54], and it is the manifestation of a well-known phenomenon: 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.²⁵ Instead, such limits *must always be recast!*

²²The limits that we call “benchmark dependent” are valid for a specific ratio of mixing angles, while the ones we call “benchmark independent” have been obtained by marginalizing over all the combinations of mixing angles allowed by the neutrino oscillation data. The latter still rely on the general properties of the model: the number of HNLs, the neutrino mass ordering and whether the HNLs behave as a Dirac-like or Majorana-like particle. As such, they are still *model dependent*.

²³The confidence limit assumes the specified mass ordering and does not take into account “priors” on the mass orderings.

²⁴This benchmark-independent limit would be much weaker for three HNLs, and non-existent for four or more HNLs, due to relaxed constraints from neutrino data.

²⁵Similarly, such limits do not apply if the HNLs have new interactions.

Figure 8. Same as figure 7, 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.

Figure 9. Original and reinterpreted exclusion limits (at $> 95\%$ CL) on the individual mixing angles $U_\alpha^2 = \sum_{I=1,2} |\Theta_{\alpha I}|^2 = x_\alpha U_{\text{tot}}^2$ and the total mixing angle $U_{\text{tot}}^2 = \sum_{\alpha=e,\mu,\tau} U_\alpha^2$ for a *Majorana-like* HNL pair and for the *normal hierarchy*. The legend is the same as in figure 7 and the rightmost figure coincides with the left panel in figure 7.

When we look at the exclusion limits obtained for the individual mixing angles (in figure 9 for the normal hierarchy and figure 10 for the inverted hierarchy), we observe that for some benchmarks the exclusion limits on individual mixing angles can sometimes be *much stronger* than the single-flavor limits. This actually reflects a rather trivial fact: if $U_\alpha^2 \ll U_\beta^2$ and the ratio $U_\alpha^2 : U_\beta^2$ is fixed, setting a limit on U_β^2 automatically sets a much stronger limit on U_α^2 (e.g. the limit set on U_e^2 for benchmark 10 in the IH indirectly sets a limit on U_μ^2 , which is enhanced by the ratio of the two mixing angles, in this case $U_\mu^2/U_e^2 \sim 1/2000$). In the same way we obtain an indirect limit (filled gray region) on

Figure 10. Same as figure 9, for a *Majorana-like* HNL pair and the *inverted hierarchy*.

Figure 11. Same as figure 9, for a *Dirac-like* HNL pair and the *normal hierarchy*. The legend is the same as in figure 8.

Figure 12. Same as figure 11, for a *Dirac-like* HNL pair and the *inverted hierarchy*.

the τ 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 τ 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 what happens when performing a scan over the parameter space in order to e.g. combine constraints from multiple sources, which may be complementary if they probe different combinations of mixing angles. For instance, we expect that future experimental results (such as excluding one neutrino mass hierarchy, or observing/setting limits on neutrinoless double-beta decay) will 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.

4.2 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 differs 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 ref. [1] 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 has therefore *no sensitivity* to Dirac-like 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 8, 11 and 12. 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 the limits for the total mixing angle (figure 8) are always weaker than the corresponding 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 the search has no sensitivity. Looking at the colored, 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 ratio of mixing angle chosen. 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 11 and 12), as well as some benchmark-independent limits (for this specific seesaw model with a Dirac-like HNL pair; see the gray filled area). 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 show a stronger dependence on their ratio than limits for Majorana-like HNLs, resulting in weaker benchmark-independent exclusion limits (filled gray 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 (from this search).

5 Conclusion & outlook

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*. As discussed in section 2, these simplified models are incompatible with the observed neutrino masses and mixing pattern. 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 [1] 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. To this end, we have implemented a simplified version of the analysis employed by ATLAS in ref. [1]. This reinterpretation was described in details in section 3.

Furthermore, as discussed in section 2.4, 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 (in their prompt HNL search), different diagrams, which depend on different combinations of mixing angles, contribute to the signal regions 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 is no sensitivity at all under the single-flavor mixing assumption! In order to handle both the Majorana- and Dirac-like cases, we have performed the reinterpretation for each of them separately. The results were respectively presented in sections 4.1 and 4.2.

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

- The exclusion limit on the total mixing angle U_{tot}^2 is always weaker (sometimes by more than 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 model 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 their ratio.

For Dirac-like HNL 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, 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_e^2 or U_μ^2 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 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.

Our results show that the reinterpretation of the exclusion limits is a necessary step in order to test HNL models which differ from those directly probed by an experiment. In particular, if one interprets the reported limits on some parameter in a given model as exclusion limits on the same parameter in a *different* model, they risk wrongly excluding part of the parameter space within the latter. This of course does not affect the validity of the limits set by the experiment for the “one-HNL, single-flavor” benchmarks; it just means that one should be cautious when investigating models other than those two initial benchmarks.

When assuming specific choices of model parameters (as in parameter scans), stronger constraints can often be derived for the individual mixing angles. In the case of two HNLs, benchmark-independent constraints can also be derived by marginalizing over all the combinations of mixing angles allowed by neutrino data. For three or more HNLs, we expect most of the above results to remain valid, with the notable exception of the marginalized limits, which become much weaker or even non-existent due to the significantly weaker constraints from neutrino data [65, 66].

For experimental results to be useful for constraining a wide range of model and parameters, it is therefore desirable to cast them into a form which allows them to be easily reinterpreted, bearing in mind that the main “drivers” for such interpretations — 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 for an easy reinterpretation of the exclusion limits.

5.2 Wish-list for a painless reinterpretation of future experimental results

The LHC collaborations typically conduct searches in terms of simplified models. Theorists, on the other hand, investigate models which address some of the shortcomings of the SM. Those are typically more complicated, and it is therefore necessary to reinterpret the search results in order to test them. In order to facilitate this reinterpretation, one would greatly benefit from the following data being reported alongside the analysis (see also the recommendations in refs. [64, 112–114]):

- The *observed bin counts*.
- The various *efficiencies* needed to evaluate the signal using the method described in sections 3.2 and 3.3, i.e.:
 - The prompt efficiency $\epsilon_{P,b}^0$ for every *process* P (as defined above, see footnote 8 on page 8) and every bin b in all signal regions. In simple cases there is a one-to-one correspondence between a *Feynman diagram* and a *process* P , as in the charge-current decays considered in this paper. Ideally, all possible processes contributing to the search signature should be included. In the present case this would mean: (a) single and mixed flavor processes; (b) LNV and LNC processes; (c) processes mediated by charge currents, neutral currents and by their interference.
 - If the parametrization in eq. (3.6) (or a modification thereof) allows reproducing the actual efficiency *even approximately*, report the relevant parameters such as the lifetime cutoff τ_0 in our case.

This slightly differs from the recommendations of the LHC Reinterpretation Forum [64], which advocates for releasing the object-level efficiencies in order to enable more general reinterpretations. Since the scope of the present reinterpretation is restricted to HNL models, those are not needed, and instead the signal can be more easily and accurately estimated using the simplified signal extrapolation method presented in sections 3.2 and 3.3.²⁶ However, we agree with their recommendation to (among many other things)

²⁶Note that our proposed reinterpretation method and the one advocated by the LHC Reinterpretation Forum are complementary: they are located in different parts of the “spectrum” of reinterpretations presented in ref. [113].

break down the efficiencies for each signal region (or bin b), each topology or final state, and each particle lifetime τ . This directly corresponds to our $\epsilon_{P,b}(\tau)$ if we include neutrinos in the final state P (which we called “process” to avoid confusion with the *visible* final state often used by experiments). As an example of how to report these per-process, per-bin efficiencies, appendix A.3 describes the JSON files containing the efficiencies computed using our simplified cut flow. A similar layout could be used to report the actual signal efficiencies from the experiment.

- For the background it is important to release *the likelihood function*. This can be either:
 - The “full” likelihood, including every background component and nuisance parameter used in the analysis (to the extent that this is possible). This can be done using tools such as `HistFactory` [112, 115] or `pyhf` [116].
 - A simplified likelihood, containing only the dominant background components and nuisance parameters (see e.g. ref. [117] or the `simplify` [118] package).
 - The covariance matrix of the background [117], for all the signal bins, across *all signal regions* (since they need to be fitted together when considering non-trivial models with e.g. both electron and muon mixing).

This is in line with the recommendations from the LHC Reinterpretation Forum [64]. Finally, to ensure that the reported likelihood is accurate enough for performing a reinterpretation, it is important to validate it, e.g. by comparing the resulting limits with those obtained using the full analysis.

To go further and to recast the analysis to a different class of models, which include Feynman diagrams not initially considered, one needs to be able to re-implement the cut flow, rather than use the efficiencies themselves. This requires knowing the efficiency maps for non-trivial cuts such as ID and isolation (as a function of both p_T and η). These maps should be *conditional* on the cuts which appear before them in the cut flow, i.e. they should be computed after applying the cuts appearing before them. This is in line with the recommendation of the LHC Reinterpretation Forum to report *analysis-specific* efficiencies [64].

Acknowledgments

JLT and OR are thankful to C. Appelt, W. Buttinger, M. Danninger, H. Lacker, D. Trischuk, M. Wielers, S. Xella for many useful discussion regarding various aspects of the ATLAS HNL analysis. JLT would like to thank F. Thiele for the many explanations regarding his Ph.D. thesis. JLT is grateful to G. Lanfranchi and M. Shaposhnikov for their constructive comments on his own Ph.D. thesis. IT is thankful to J. Klarić and K. Urquía for helpful comments on the manuscript. 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 and JLT has been supported by ERC-AdG-2015 grant 694896 and by the Swiss National Science Foundation Excellence grant 200020B_182864.

A Ancillary files

In order to simplify the interpretation of experimental results within realistic HNL models, we are including a number of data files along with the present publication. They can be used to generate the relevant signal samples, or to implement the extrapolation method presented in section 3.2. These files can be found in the companion Zenodo record [119].

A.1 Card files for the Monte-Carlo event generation

The `/attachments/card_files` folder contains the MADGRAPH card files (ending in `.dat`) and scripts (ending in `.txt`) for generating the signal samples used in this analysis, as well as for computing the total HNL width. Due to the OSSF veto, only processes with no opposite-charge same-flavor lepton pairs have been included. Additional relevant processes can easily be added by modifying the `generate` and `add process` lines in the `*.txt` files. All samples (except the ones used to compute the HNL width, which are generated at parton level) are generated at leading order, include up to two hard jets, and are showered and hadronized using PYTHIA 8. This is essential for obtaining a realistic W spectrum. The shower parameters could probably benefit from further tuning, and further improvements in the W spectrum accuracy are expected at NLO (using a suitable model). To allow computing the signal efficiencies, all cuts have been disabled in the run card (with the exception of the maximum $|\eta_{\text{jet}}|$ which needs to be set to 5 for correct matching).

A.2 Signal cross sections

The cross sections for the various processes considered in this analysis, as well as the total HNL width (both computed using MADGRAPH as described in section 3.2), are provided as JSON files in the `/attachments/cross_sections` folder.

The file `total_hnl_width.json` contains the total HNL width $\hat{\Gamma}_\alpha(M_N)$ (expressed in GeV), computed for the 5 mass points used in this analysis, and under the assumption of unit mixing with a single flavor α , for each flavor. The total HNL width can then be computed for any combinations of mixing angles using eq. (3.2). The file is organized as two nested dictionaries, with the first key denoting the HNL mass M_N , and the second one the flavor α for which the total width $\hat{\Gamma}_\alpha(M_N)$ has been computed for a unit mixing angle $|\Theta_\alpha|^2 = 1$ (with `Wtot_e` for $\alpha = e$, `Wtot_mu` for μ and `Wtot_tau` for τ).

The file `cross_sections.json` contains the reference cross sections σ_P^{ref} (in pb) for all the processes P considered in this analysis, expressed for $|\Theta|_{\text{ref}}^2 = 1$ and $\Gamma_{\text{ref}} = 10^{-5}$ GeV. The file is organized as two nested dictionaries, with the first key denoting the HNL mass M_N and the second the process P . The correspondence between the key and the physical process can be found in table 7.

A.3 Signal efficiencies

The efficiencies resulting from the event selection described in section 3.1, as well as their parametrization according to eq. (3.6) (as discussed in section 3.3) can respectively be found in the files `efficiencies.json` and `fitted_efficiencies.json` in the `/attachments/efficiencies` folder.

Key	Process
lnc_e+mu-e+	$W^+ \rightarrow e^+(N \rightarrow \mu^- e^+ \nu_e)$
lnc_e-mu+e-	$W^- \rightarrow e^-(N \rightarrow \mu^+ e^- \bar{\nu}_e)$
lnc_mu+e-mu+	$W^+ \rightarrow \mu^+(N \rightarrow e^- \mu^+ \nu_\mu)$
lnc_mu-e+mu-	$W^- \rightarrow \mu^-(N \rightarrow e^+ \mu^- \bar{\nu}_\mu)$
lnv_e+e+mu-	$W^+ \rightarrow e^+(N \rightarrow e^+ \mu^- \bar{\nu}_\mu)$
lnv_e-e-mu+	$W^- \rightarrow e^-(N \rightarrow e^- \mu^+ \nu_\mu)$
lnv_mu+mu+e-	$W^+ \rightarrow \mu^+(N \rightarrow \mu^+ e^- \bar{\nu}_e)$
lnv_mu-mu-e+	$W^- \rightarrow \mu^-(N \rightarrow \mu^- e^+ \nu_e)$

Table 7. Correspondence between the process key and the actual process.

The file `efficiencies.json` is organized as follows. The data is located in a triply nested dictionary under the `data` key: the first level corresponds to the HNL mass hypothesis M_N , the second to the process key (cf. table 7) and the third to the $M(l_{\text{sublead}}, l')$ bin for which the efficiency is computed. The values of the bottom-most dictionary are lists containing the efficiencies for a number of HNL lifetimes, as listed in meters in `levels/lifetime`.

Finally, the file `fitted_efficiencies.json` is also organized as a triply nested dictionary, with the first level corresponding to the HNL mass M_N , the second to the process key, and where the third level denotes the fit parameter from eq. (3.6). `tau0` is for τ_0 , `epsilon0_total` for ϵ_0 (the unbinned prompt efficiency), and `epsilon0_binned` is a list containing the prompt efficiencies $\epsilon_{0,b}$ for the five $M(l_{\text{sublead}}, l')$ bins b (in the same order as in `efficiencies.json`). The layout described here (or a similar one) can be used by experiments to report their signal efficiencies in a way that ensures that theorists will be able to compute the expected signal for arbitrary choices of mixing angles.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] 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](https://arxiv.org/abs/1905.09787)] [[INSPIRE](#)].
- [2] M. Shaposhnikov, *Is there a new physics between electroweak and Planck scales?*, in *Astroparticle Physics: Current Issues, 2007 (APCI07)*, (2007) [[arXiv:0708.3550](https://arxiv.org/abs/0708.3550)] [[INSPIRE](#)].
- [3] S. Alekhin et al., *A facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case*, *Rept. Prog. Phys.* **79** (2016) 124201 [[arXiv:1504.04855](https://arxiv.org/abs/1504.04855)] [[INSPIRE](#)].
- [4] J. Beacham et al., *Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report*, *J. Phys. G* **47** (2020) 010501 [[arXiv:1901.09966](https://arxiv.org/abs/1901.09966)] [[INSPIRE](#)].

- [5] R.K. Ellis et al., *Physics Briefing Book: Input for the European Strategy for Particle Physics Update 2020*, [arXiv:1910.11775](#) [INSPIRE].
- [6] T. Asaka, S. Blanchet and M. Shaposhnikov, *The ν MSM, dark matter and neutrino masses*, *Phys. Lett. B* **631** (2005) 151 [[hep-ph/0503065](#)] [INSPIRE].
- [7] T. Asaka and M. Shaposhnikov, *The ν MSM, dark matter and baryon asymmetry of the universe*, *Phys. Lett. B* **620** (2005) 17 [[hep-ph/0505013](#)] [INSPIRE].
- [8] A. Boyarsky, O. Ruchayskiy and M. Shaposhnikov, *The Role of sterile neutrinos in cosmology and astrophysics*, *Ann. Rev. Nucl. Part. Sci.* **59** (2009) 191 [[arXiv:0901.0011](#)] [INSPIRE].
- [9] M. Drewes, *The Phenomenology of Right Handed Neutrinos*, *Int. J. Mod. Phys. E* **22** (2013) 1330019 [[arXiv:1303.6912](#)] [INSPIRE].
- [10] F.F. Deppisch, P.S. Bhupal Dev and A. Pilaftsis, *Neutrinos and Collider Physics*, *New J. Phys.* **17** (2015) 075019 [[arXiv:1502.06541](#)] [INSPIRE].
- [11] J.C. Pati and A. Salam, *Lepton Number as the Fourth Color*, *Phys. Rev. D* **10** (1974) 275 [Erratum *ibid.* **11** (1975) 703] [INSPIRE].
- [12] R.N. Mohapatra and J.C. Pati, *A Natural Left-Right Symmetry*, *Phys. Rev. D* **11** (1975) 2558 [INSPIRE].
- [13] R.N. Mohapatra and J.C. Pati, *Left-Right Gauge Symmetry and an Isoconjugate Model of CP-violation*, *Phys. Rev. D* **11** (1975) 566 [INSPIRE].
- [14] G. Senjanović and R.N. Mohapatra, *Exact Left-Right Symmetry and Spontaneous Violation of Parity*, *Phys. Rev. D* **12** (1975) 1502 [INSPIRE].
- [15] P. Minkowski, *$\mu \rightarrow e\gamma$ at a Rate of One Out of 10^9 Muon Decays?*, *Phys. Lett. B* **67** (1977) 421 [INSPIRE].
- [16] T. Yanagida, *Horizontal gauge symmetry and masses of neutrinos*, *Conf. Proc. C* **7902131** (1979) 95 [INSPIRE].
- [17] S.L. Glashow, *The Future of Elementary Particle Physics*, *NATO Sci. Ser. B* **61** (1980) 687 [INSPIRE].
- [18] M. Gell-Mann, P. Ramond and R. Slansky, *Complex Spinors and Unified Theories*, *Conf. Proc. C* **790927** (1979) 315 [[arXiv:1306.4669](#)] [INSPIRE].
- [19] R.N. Mohapatra and G. Senjanović, *Neutrino Mass and Spontaneous Parity Nonconservation*, *Phys. Rev. Lett.* **44** (1980) 912 [INSPIRE].
- [20] R.N. Mohapatra and G. Senjanović, *Neutrino Masses and Mixings in Gauge Models with Spontaneous Parity Violation*, *Phys. Rev. D* **23** (1981) 165 [INSPIRE].
- [21] J. Schechter and J.W.F. Valle, *Neutrino Masses in $SU(2) \times U(1)$ Theories*, *Phys. Rev. D* **22** (1980) 2227 [INSPIRE].
- [22] J. Schechter and J.W.F. Valle, *Neutrino Decay and Spontaneous Violation of Lepton Number*, *Phys. Rev. D* **25** (1982) 774 [INSPIRE].
- [23] M. Fukugita and T. Yanagida, *Resurrection of grand unified theory baryogenesis*, *Phys. Rev. Lett.* **89** (2002) 131602 [[hep-ph/0203194](#)] [INSPIRE].
- [24] S. Davidson, E. Nardi and Y. Nir, *Leptogenesis*, *Phys. Rept.* **466** (2008) 105 [[arXiv:0802.2962](#)] [INSPIRE].

- [25] M. Shaposhnikov, *Baryogenesis*, *J. Phys. Conf. Ser.* **171** (2009) 012005 [[INSPIRE](#)].
- [26] E.K. Akhmedov, V.A. Rubakov and A.Y. Smirnov, *Baryogenesis via neutrino oscillations*, *Phys. Rev. Lett.* **81** (1998) 1359 [[hep-ph/9803255](#)] [[INSPIRE](#)].
- [27] A. Pilaftsis, *The Little Review on Leptogenesis*, *J. Phys. Conf. Ser.* **171** (2009) 012017 [[arXiv:0904.1182](#)] [[INSPIRE](#)].
- [28] M. Shaposhnikov, *The ν MSM, leptonic asymmetries, and properties of singlet fermions*, *JHEP* **08** (2008) 008 [[arXiv:0804.4542](#)] [[INSPIRE](#)].
- [29] L. Canetti and M. Shaposhnikov, *Baryon Asymmetry of the Universe in the ν MSM*, *JCAP* **09** (2010) 001 [[arXiv:1006.0133](#)] [[INSPIRE](#)].
- [30] M. Drewes et al., *ARS Leptogenesis*, *Int. J. Mod. Phys. A* **33** (2018) 1842002 [[arXiv:1711.02862](#)] [[INSPIRE](#)].
- [31] J. Klarić, M. Shaposhnikov and I. Timiryasov, *Uniting Low-Scale Leptogenesis Mechanisms*, *Phys. Rev. Lett.* **127** (2021) 111802 [[arXiv:2008.13771](#)] [[INSPIRE](#)].
- [32] J. Klarić, M. Shaposhnikov and I. Timiryasov, *Reconciling resonant leptogenesis and baryogenesis via neutrino oscillations*, *Phys. Rev. D* **104** (2021) 055010 [[arXiv:2103.16545](#)] [[INSPIRE](#)].
- [33] A. Boyarsky, M. Drewes, T. Lasserre, S. Mertens and O. Ruchayskiy, *Sterile neutrino Dark Matter*, *Prog. Part. Nucl. Phys.* **104** (2019) 1 [[arXiv:1807.07938](#)] [[INSPIRE](#)].
- [34] J. Ghiglieri and M. Laine, *Sterile neutrino dark matter via coinciding resonances*, *JCAP* **07** (2020) 012 [[arXiv:2004.10766](#)] [[INSPIRE](#)].
- [35] R.E. Shrock, *New Tests For, and Bounds On, Neutrino Masses and Lepton Mixing*, *Phys. Lett. B* **96** (1980) 159 [[INSPIRE](#)].
- [36] R.E. Shrock, *General Theory of Weak Leptonic and Semileptonic Decays. I. Leptonic Pseudoscalar Meson Decays, with Associated Tests For, and Bounds on, Neutrino Masses and Lepton Mixing*, *Phys. Rev. D* **24** (1981) 1232 [[INSPIRE](#)].
- [37] R.E. Shrock, *General Theory of Weak Processes Involving Neutrinos. II. Pure Leptonic Decays*, *Phys. Rev. D* **24** (1981) 1275 [[INSPIRE](#)].
- [38] M. Gronau, C.N. Leung and J.L. Rosner, *Extending Limits on Neutral Heavy Leptons*, *Phys. Rev. D* **29** (1984) 2539 [[INSPIRE](#)].
- [39] D. Gorbunov and M. Shaposhnikov, *How to find neutral leptons of the ν MSM?*, *JHEP* **10** (2007) 015 [*Erratum ibid.* **11** (2013) 101] [[arXiv:0705.1729](#)] [[INSPIRE](#)].
- [40] A. Atre, T. Han, S. Pascoli and B. Zhang, *The Search for Heavy Majorana Neutrinos*, *JHEP* **05** (2009) 030 [[arXiv:0901.3589](#)] [[INSPIRE](#)].
- [41] 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](#)] [[INSPIRE](#)].
- [42] BELLE collaboration, *Search for heavy neutrinos at Belle*, *Phys. Rev. D* **87** (2013) 071102 [*Erratum ibid.* **95** (2017) 099903] [[arXiv:1301.1105](#)] [[INSPIRE](#)].
- [43] LHCb collaboration, *Search for Majorana neutrinos in $B^- \rightarrow \pi^+ \mu^- \mu^-$ decays*, *Phys. Rev. Lett.* **112** (2014) 131802 [[arXiv:1401.5361](#)] [[INSPIRE](#)].
- [44] E949 collaboration, *Search for heavy neutrinos in $K^+ \rightarrow \mu^+ \nu_H$ decays*, *Phys. Rev. D* **91** (2015) 052001 [*Erratum ibid.* **91** (2015) 059903] [[arXiv:1411.3963](#)] [[INSPIRE](#)].

- [45] 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](#)] [[INSPIRE](#)].
- [46] CMS collaboration, *Search for heavy Majorana neutrinos in $\mu^\pm\mu^\pm +$ jets events in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *Phys. Lett. B* **748** (2015) 144 [[arXiv:1501.05566](#)] [[INSPIRE](#)].
- [47] 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. D* **97** (2018) 015023 [[arXiv:1708.09395](#)] [[INSPIRE](#)].
- [48] NA62 collaboration, *Search for heavy neutral lepton production in K^+ decays*, *Phys. Lett. B* **778** (2018) 137 [[arXiv:1712.00297](#)] [[INSPIRE](#)].
- [49] SHiP collaboration, *Prospects of the SHiP and NA62 experiments at CERN for hidden sector searches*, *PoS NuFact2017* (2017) 139 [[arXiv:1712.01768](#)] [[INSPIRE](#)].
- [50] A. Izmaylov and S. Suvorov, *Search for heavy neutrinos in the ND280 near detector of the T2K experiment*, *Phys. Part. Nucl.* **48** (2017) 984 [[INSPIRE](#)].
- [51] 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** (2018) 221801 [[arXiv:1802.02965](#)] [[INSPIRE](#)].
- [52] NA62 collaboration, *Search for heavy neutral lepton production in K^+ decays to positrons*, *Phys. Lett. B* **807** (2020) 135599 [[arXiv:2005.09575](#)] [[INSPIRE](#)].
- [53] NA62 collaboration, *Search for K^+ decays to a muon and invisible particles*, *Phys. Lett. B* **816** (2021) 136259 [[arXiv:2101.12304](#)] [[INSPIRE](#)].
- [54] A. Abada, N. Bernal, M. Losada and X. Marcano, *Inclusive Displaced Vertex Searches for Heavy Neutral Leptons at the LHC*, *JHEP* **01** (2019) 093 [[arXiv:1807.10024](#)] [[INSPIRE](#)].
- [55] SHiP collaboration, *Sensitivity of the SHiP experiment to Heavy Neutral Leptons*, *JHEP* **04** (2019) 077 [[arXiv:1811.00930](#)] [[INSPIRE](#)].
- [56] K. Cranmer and I. Yavin, *RECAST: Extending the Impact of Existing Analyses*, *JHEP* **04** (2011) 038 [[arXiv:1010.2506](#)] [[INSPIRE](#)].
- [57] A. Buckley et al., *Rivet user manual*, *Comput. Phys. Commun.* **184** (2013) 2803 [[arXiv:1003.0694](#)] [[INSPIRE](#)].
- [58] E. Conte, B. Fuks and G. Serret, *MadAnalysis 5, A User-Friendly Framework for Collider Phenomenology*, *Comput. Phys. Commun.* **184** (2013) 222 [[arXiv:1206.1599](#)] [[INSPIRE](#)].
- [59] E. Conte and B. Fuks, *MadAnalysis 5: status and new developments*, *J. Phys. Conf. Ser.* **523** (2014) 012032 [[arXiv:1309.7831](#)] [[INSPIRE](#)].
- [60] 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 [[arXiv:1312.2591](#)] [[INSPIRE](#)].
- [61] S. Kraml et al., *SModelS: a tool for interpreting simplified-model results from the LHC and its application to supersymmetry*, *Eur. Phys. J. C* **74** (2014) 2868 [[arXiv:1312.4175](#)] [[INSPIRE](#)].
- [62] M. Papucci, K. Sakurai, A. Weiler and L. Zeune, *Fastlim: a fast LHC limit calculator*, *Eur. Phys. J. C* **74** (2014) 3163 [[arXiv:1402.0492](#)] [[INSPIRE](#)].

- [63] D. Barducci et al., *Framework for Model Independent Analyses of Multiple Extra Quark Scenarios*, *JHEP* **12** (2014) 080 [[arXiv:1405.0737](#)] [[INSPIRE](#)].
- [64] LHC REINTERPRETATION FORUM collaboration, *Reinterpretation of LHC Results for New Physics: Status and Recommendations after Run 2*, *SciPost Phys.* **9** (2020) 022 [[arXiv:2003.07868](#)] [[INSPIRE](#)].
- [65] 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](#)] [[INSPIRE](#)].
- [66] 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. C* **80** (2020) 569 [[arXiv:1908.02302](#)] [[INSPIRE](#)].
- [67] J.-L. Tastet and I. Timiryasov, *Dirac vs. Majorana HNLs (and their oscillations) at SHiP*, *JHEP* **04** (2020) 005 [[arXiv:1912.05520](#)] [[INSPIRE](#)].
- [68] D. Boyanovsky, *Nearly degenerate heavy sterile neutrinos in cascade decay: mixing and oscillations*, *Phys. Rev. D* **90** (2014) 105024 [[arXiv:1409.4265](#)] [[INSPIRE](#)].
- [69] G. Cvetič, C.S. Kim, R. Kogler and J. Zamora-Saá, *Oscillation of heavy sterile neutrino in decay of $B \rightarrow \mu e \pi$* , *Phys. Rev. D* **92** (2015) 013015 [[arXiv:1505.04749](#)] [[INSPIRE](#)].
- [70] G. Anamiati, M. Hirsch and E. Nardi, *Quasi-Dirac neutrinos at the LHC*, *JHEP* **10** (2016) 010 [[arXiv:1607.05641](#)] [[INSPIRE](#)].
- [71] S. Antusch, E. Cazzato and O. Fischer, *Resolvable heavy neutrino-antineutrino oscillations at colliders*, *Mod. Phys. Lett. A* **34** (2019) 1950061 [[arXiv:1709.03797](#)] [[INSPIRE](#)].
- [72] 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. D* **97** (2018) 015018 [[arXiv:1709.06553](#)] [[INSPIRE](#)].
- [73] 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](#)] [[INSPIRE](#)].
- [74] P. Hernández, J. Jones-Pérez and O. Suarez-Navarro, *Majorana vs Pseudo-Dirac Neutrinos at the ILC*, *Eur. Phys. J. C* **79** (2019) 220 [[arXiv:1810.07210](#)] [[INSPIRE](#)].
- [75] M. Shaposhnikov, *A Possible symmetry of the ν MSM*, *Nucl. Phys. B* **763** (2007) 49 [[hep-ph/0605047](#)] [[INSPIRE](#)].
- [76] J. Kersten and A.Y. Smirnov, *Right-Handed Neutrinos at CERN LHC and the Mechanism of Neutrino Mass Generation*, *Phys. Rev. D* **76** (2007) 073005 [[arXiv:0705.3221](#)] [[INSPIRE](#)].
- [77] M. Drewes, J. Klarić and P. Klose, *On lepton number violation in heavy neutrino decays at colliders*, *JHEP* **11** (2019) 032 [[arXiv:1907.13034](#)] [[INSPIRE](#)].
- [78] K. Bondarenko et al., *An allowed window for heavy neutral leptons below the kaon mass*, *JHEP* **07** (2021) 193 [[arXiv:2101.09255](#)] [[INSPIRE](#)].
- [79] PARTICLE DATA GROUP collaboration, *Review of Particle Physics*, *PTEP* **2020** (2020) 083C01 [[INSPIRE](#)].
- [80] S. Davidson, G. Isidori and A. Strumia, *The Smallest neutrino mass*, *Phys. Lett. B* **646** (2007) 100 [[hep-ph/0611389](#)] [[INSPIRE](#)].

- [81] S.M. Bilenky, *Neutrino in Standard Model and beyond*, *Phys. Part. Nucl.* **46** (2015) 475 [[arXiv:1501.00232](#)] [[INSPIRE](#)].
- [82] 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](#)] [[INSPIRE](#)].
- [83] APPEC COMMITTEE collaboration, *Double Beta Decay APPEC Committee Report*, [arXiv:1910.04688](#) [[INSPIRE](#)].
- [84] T2K collaboration, *Constraint on the matter-antimatter symmetry-violating phase in neutrino oscillations*, *Nature* **580** (2020) 339 [Erratum *ibid.* **583** (2020) E16] [[arXiv:1910.03887](#)] [[INSPIRE](#)].
- [85] DUNE collaboration, *The DUNE Far Detector Interim Design Report Volume 1: Physics, Technology and Strategies*, [arXiv:1807.10334](#) [[INSPIRE](#)].
- [86] J.A. Casas and A. Ibarra, *Oscillating neutrinos and $\mu \rightarrow e, \gamma$* , *Nucl. Phys. B* **618** (2001) 171 [[hep-ph/0103065](#)] [[INSPIRE](#)].
- [87] A. Abada, S. Davidson, A. Ibarra, F.X. Josse-Michaux, M. Losada and A. Riotto, *Flavour Matters in Leptogenesis*, *JHEP* **09** (2006) 010 [[hep-ph/0605281](#)] [[INSPIRE](#)].
- [88] S. Eijima, M. Shaposhnikov and I. Timiryasov, *Parameter space of baryogenesis in the ν MSM*, *JHEP* **07** (2019) 077 [[arXiv:1808.10833](#)] [[INSPIRE](#)].
- [89] M. Drewes, B. Garbrecht, D. Gueter and J. Klaric, *Testing the low scale seesaw and leptogenesis*, *JHEP* **08** (2017) 018 [[arXiv:1609.09069](#)] [[INSPIRE](#)].
- [90] A. Caputo, P. Hernández, J. López-Pavon and J. Salvado, *The seesaw portal in testable models of neutrino masses*, *JHEP* **06** (2017) 112 [[arXiv:1704.08721](#)] [[INSPIRE](#)].
- [91] 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](#)] [[INSPIRE](#)].
- [92] 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](#)] [[INSPIRE](#)].
- [93] I. Esteban, M.C. Gonzalez-Garcia, M. Maltoni, T. Schwetz and A. Zhou, *NuFIT 5.0* (2020) [<http://www.nu-fit.org>].
- [94] SUPER-KAMIOKANDE collaboration, *Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV*, *Phys. Rev. D* **97** (2018) 072001 [[arXiv:1710.09126](#)] [[INSPIRE](#)].
- [95] D. Wyler and L. Wolfenstein, *Massless Neutrinos in Left-Right Symmetric Models*, *Nucl. Phys. B* **218** (1983) 205 [[INSPIRE](#)].
- [96] R.N. Mohapatra and J.W.F. Valle, *Neutrino Mass and Baryon Number Nonconservation in Superstring Models*, *Phys. Rev. D* **34** (1986) 1642 [[INSPIRE](#)].
- [97] G.C. Branco, W. Grimus and L. Lavoura, *The Seesaw Mechanism in the Presence of a Conserved Lepton Number*, *Nucl. Phys. B* **312** (1989) 492 [[INSPIRE](#)].
- [98] M.C. Gonzalez-Garcia and J.W.F. Valle, *Fast Decaying Neutrinos and Observable Flavor Violation in a New Class of Majoron Models*, *Phys. Lett. B* **216** (1989) 360 [[INSPIRE](#)].

- [99] 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](#)] [[INSPIRE](#)].
- [100] A. Roy and M. Shaposhnikov, *Resonant production of the sterile neutrino dark matter and fine-tunings in the $[\nu]$ MSM*, *Phys. Rev. D* **82** (2010) 056014 [[arXiv:1006.4008](#)] [[INSPIRE](#)].
- [101] M.B. Gavela, T. Hambye, D. Hernández and P. Hernández, *Minimal Flavour Seesaw Models*, *JHEP* **09** (2009) 038 [[arXiv:0906.1461](#)] [[INSPIRE](#)].
- [102] L. Canetti, M. Drewes, T. Frossard and M. Shaposhnikov, *Dark Matter, Baryogenesis and Neutrino Oscillations from Right Handed Neutrinos*, *Phys. Rev. D* **87** (2013) 093006 [[arXiv:1208.4607](#)] [[INSPIRE](#)].
- [103] F. Thiele, *An ATLAS Search for Sterile Neutrinos*, Ph.D. Thesis, Niels Bohr Institutet, Copenhagen, Denmark (2019) [[INSPIRE](#)].
- [104] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[arXiv:1405.0301](#)] [[INSPIRE](#)].
- [105] D. Alva, T. Han and R. Ruiz, *Heavy Majorana neutrinos from $W\gamma$ fusion at hadron colliders*, *JHEP* **02** (2015) 072 [[arXiv:1411.7305](#)] [[INSPIRE](#)].
- [106] C. Degrande, O. Mattelaer, R. Ruiz and J. Turner, *Fully-Automated Precision Predictions for Heavy Neutrino Production Mechanisms at Hadron Colliders*, *Phys. Rev. D* **94** (2016) 053002 [[arXiv:1602.06957](#)] [[INSPIRE](#)].
- [107] ATLAS collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*, *2019 JINST* **14** P12006 [[arXiv:1908.00005](#)] [[INSPIRE](#)].
- [108] ATLAS collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 578 [[arXiv:2012.00578](#)] [[INSPIRE](#)].
- [109] ATLAS collaboration, *Expected b -tagging performance with the upgraded ATLAS Inner Tracker detector at the High-Luminosity LHC*, Tech. Rep. *ATL-PHYS-PUB-2020-005*, CERN, Geneva (Mar, 2020).
- [110] R. Ruiz, *Quantitative study on helicity inversion in Majorana neutrino decays at the LHC*, *Phys. Rev. D* **103** (2021) 015022 [[arXiv:2008.01092](#)] [[INSPIRE](#)].
- [111] A.L. Read, *Presentation of search results: The CL_s technique*, *J. Phys. G* **28** (2002) 2693 [[INSPIRE](#)].
- [112] ROOT collaboration, *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*, Tech. Rep. *CERN-OPEN-2012-016*, New York U., New York (Jan, 2012).
- [113] N. Wardle, *Can we really “Re”-interpret data from the LHC?* in *PHYSTAT seminar* (2021).
- [114] C. Bierlich et al., *Robust Independent Validation of Experiment and Theory: Rivet version 3*, *SciPost Phys.* **8** (2020) 026 [[arXiv:1912.05451](#)] [[INSPIRE](#)].
- [115] ATLAS collaboration, *Reproducing searches for new physics with the ATLAS experiment through publication of full statistical likelihoods*, Tech. Rep. *ATL-PHYS-PUB-2019-029*, CERN, Geneva (Aug, 2019).
- [116] L. Heinrich, M. Feickert, G. Stark and K. Cranmer, *pyhf: pure-Python implementation of HistFactory statistical models*, *J. Open Source Softw.* **6** (2021) 2823.

- [117] CMS collaboration, *Simplified likelihood for the re-interpretation of public CMS results*, Tech. Rep. [CMS-NOTE-2017-001](#), [CERN-CMS-NOTE-2017-001](#), CERN, Geneva (Jan, 2017).
- [118] E. Schanet, *simplify*, <https://github.com/eschanet/simplify>.
- [119] Tastet, Jean-Loup, Ruchayskiy, Oleg, Timiryasov and Inar, *Ancillary files for Reinterpreting the ATLAS bounds on heavy neutral leptons in a realistic neutrino oscillation model [arXiv: 2107.12980]*, [Zenodo \[data set\]](#) **v1.0** (2021).