


Modeling heavy neutral leptons in accelerator beamlines

Komninos-John Plows^{1,*} and Xianguo Lu^{2,†}

¹University of Oxford, Department of Physics, Oxford, OX1 3PJ United Kingdom

²University of Warwick, Department of Physics, Coventry, CV4 7AL United Kingdom

 (Received 21 December 2022; accepted 16 February 2023; published 6 March 2023)

Heavy neutral leptons (HNLs) with masses $\mathcal{O}(0.1\text{--}1\text{ GeV}/c^2)$ are promising candidates for the simultaneous explanation of the smallness of the observed neutrino masses as well as the matter-antimatter asymmetry in the observable Universe. These particles can be produced in the decay of hadrons typically produced in a neutrino beamline used for oscillation experiments, and have sufficient lifetime to propagate to a near detector, where they decay to observable particles. For the approximation of a single new mass eigenstate mixing with the Standard Model via the lepton mixing matrix, a simulation framework based on the GENIE event generator has been developed. This module is designed to facilitate searches for HNL through a unified, minimal interface employing a detailed treatment of the kinematics and dynamics of massive unstable neutrinos, with a transparently organized suite of physics effects tracking the HNL from its production to its decay. These mechanisms are expounded on in the current work, underlining the rich landscape for novel, nontrivial physics that has already been identified in previous literature. This framework is an ongoing effort to provide a consistent and comprehensive description of heavy neutrinos from particle decays. We highlight use cases and future applications of interest to the accelerator neutrino community.

DOI: [10.1103/PhysRevD.107.055003](https://doi.org/10.1103/PhysRevD.107.055003)

I. INTRODUCTION

Neutrino oscillations [1,2] were proposed early in the history of neutrino physics (see, for example, [3] and references therein). However, it was the discovery of the deficit in solar neutrino events in the Davis experiment [4] that first gave indications for physics beyond the Standard Model. The discovery of neutrino flavor conversion first in solar [5] and atmospheric [6] neutrinos, and then in reactor antineutrinos [7], confirmed the hypothesis that neutrinos oscillate between flavor eigenstates [8], which runs contrary to the Standard Model expectation that all neutrinos, being left-chiral, are massless. The oscillation parameters (mass splittings, mixing angles, and CP violating phase δ_{CP}) have since become the subject of precise measurement campaigns from a plethora of detectors from different baselines past [5,6,9–20] and present [7,21–27]. Further experiments are being designed to increase the precision to which the oscillation parameters are known [28–31], and a comprehensive effort to produce global fits to extant data is underway [8].

At the same time, the masses of the active neutrinos, albeit nonzero, are very small. Direct measurements using beta decay at KATRIN ([32]) currently yield a bound on the “electron neutrino mass” of $(\sum_i |U_{ei}|^2 m_{\nu_i}^2)^{1/2} \lesssim 0.8\text{ eV}$, where U is the 3×3 PMNS matrix and i runs over the three known mass eigenstates. At the same time, indirect bounds on the sum of the masses $\sum_i m_{\nu_i}$ within a 7-parameter cosmological model ($\Lambda\text{CDM} + \sum_i m_{\nu_i}$) are even more stringent. Constraints on the cosmic microwave background from Planck 2018 data are currently at $\sum_i m_{\nu_i} < 0.26\text{ eV}$, dropping to $< 0.13\text{ eV}$ when combined with measurements of the scale of baryon acoustic oscillations [33]. Indirect bounds on neutrino mass via the “effective mass” $|m_{\beta\beta}| = |\sum_i U_{ei}^2 m_{\nu_i}|$ may also be obtained from neutrinoless double-beta decay experiments [34], under the assumption that neutrinoless double-beta decay is mediated by active Majorana neutrinos only.

To explain the nonzero albeit small observed neutrino masses, a seesaw mechanism of mass generation is usually invoked [35–37]. Seesaw mechanisms posit new neutrino degrees of freedom, which carry no charge under electro-weak symmetry $SU(2)_L \times U(1)_Y$: that is, the new neutrino states are sterile. Notably, both Dirac and Majorana fields allow for viable seesaws. The mass eigenstates associated with these new fields, whose scale is dependent on the details of the seesaw mechanism, are termed heavy neutral leptons, or HNL. These mix with the light neutrino mass states into the observed flavor eigenstates $\nu_{e,\mu,\tau}$, and

*komninos-john.plows@mansfield.ox.ac.uk

†xianguo.lu@warwick.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Funded by SCOAP³.

therefore can enter into interactions involving visible Standard Model particles. The smallness of the active neutrino masses implies, for a wide variety of seesaw models [38–40], that either the HNL Yukawa couplings are very small, or that their masses are much larger than currently accessible experimental scales. However, so-called “low-scale seesaw” models allow for HNL with masses in the range $\mathcal{O}(100 \text{ MeV}/c^2 - 1 \text{ TeV}/c^2)$ [41–46].

Apart from the ability to explain the existence and smallness of observed neutrino masses, HNL can also mediate the observed matter-antimatter asymmetry in the Universe [47,48]. Moreover, there are scenarios ([38,49]) in which an HNL of keV-scale mass can act as a warm-dark-matter component, explaining at least part of the dark matter content in the Universe. In fact, a detected unidentified emission line at $E_\gamma \simeq 3.55 \text{ keV}$ in stacked galactic spectra obtained with the *XMM-Newton* instrument [50] had been proposed to result from the decay of a $7.1 \text{ keV}/c^2$ mass HNL, but subsequent searches for the line in different astronomical features have not found significant evidence for an excess at 3.55 keV [51,52]. This realization has motivated searches for HNL at accelerator facilities by means of searches for HNL decays [53–60], searches for anomalous peaks in the spectra of leptons from meson decays [61–66] (see [67] for an overview), as well as significant interest for searches in future experiments [68–77]. In addition to these searches, a series of precision measurements of the SM predictions for the decays of known particles, such as μ , τ , π , K , can be used to derive bounds on the HNL parameter space (see [78] for an overview).

In the case of decay searches, it is crucial to accurately model the flux of HNL at the detector, as well as maintain precise control on the backgrounds to detection channels. Standard Model neutrino measurements make use of a suite of “standard” neutrino event generators. These support wide-ranging analyses from sensitivity studies all the way through to cross section measurements to oscillation analyses (see [79] for an overview). This is not the case with new-physics signatures, only a limited selection of which is implemented in the tools currently available to the community [80]. With the exception of the *ACHILLES* [81] and *DarkNews* [82] generators, as well as of a partial implementation in the *GENIE* [83,84] generator, there have been few efforts to incorporate HNL physics into event generators. The *GENIE* neutrino event generator is an event generator that is used by existing and future experiments. It leverages a *ROOT*-based [85,86] *C++* framework, including a series of mature flux and geometry drivers [87], to take advantage of object-oriented techniques in the creation of neutrino events. *GENIE* currently offers an implementation of short-lived HNL produced in Standard Model neutrino upscattering from interactions with the nucleus [88], motivated by [89], that, however, does not model long-lived HNL produced in the beamline.

For searches individually implementing HNL physics (see, e.g., [53,57]), additional efforts to interface with the beamline simulation and implement the detector geometry are needed. Additionally, HNL simulations must correctly describe the production kinematics of the HNL taking into account their non-negligible mass, as well as the propagation and decay of these particles. Since HNL are unstable, this implies properly implementing the probability of decay inside (or immediately upstream) of the detector, as well as describing the angular dependence of the decay products induced by the polarization of the HNL.

In this work, we not only take care of the beamline interface and detector integration in a fully controllable manner, but also investigate the kinematics and propagation/decay of HNL in detail. Our long-lived HNL implementation starts from a general description of “parent” particles (such as pseudoscalar mesons) traveling in a beamline. It is a general framework that constructs HNL from the decay of these parents, calculates the decay rates to all kinematically accessible final states, and returns the decay products of HNL decays to those final states that the user wants to simulate, inside a general detector volume the user provides as input. The choice of which channels are desirable, as well as the beamline simulation and the detector description, is fully configurable. This has the advantage of factorizing out the complexity of the beamline simulation and detector geometry, focusing instead on the core physics of massive neutrinos while simultaneously providing output that is usable for a wide array of beamlines and detector setups. This allows for the use of this module for optimization studies in a broad range of parameters of interest. This framework takes care of all kinematical calculations, and it is our goal to support present and future HNL models via a replaceable module defined by the HNL production and decay rates. We implement this framework in *GENIE*; a quick guide to running the *GENIE* implementation can be seen in Appendix A 1.

The framework we introduce not only handles the various physics effects of massive neutrinos in a general way, but also produces standard *GENIE* output, enabling experiments to directly use its results in their Monte Carlo production chain.

This paper is organized as follows. In Secs. II and III, we will detail the specific model treated in this implementation, and present a brief overview of the HNL simulation chains in previous searches. The model can be decomposed into the physical production mechanism from hadron decays following proton-target interactions, and the Lagrangian from which the production and decay rates into the kinematically accessible channels are calculated. In Sec. IV, we will describe in detail the implementation of HNL production, propagation, and decay in the *BeamHNL* module of *GENIE* v3. Next, in Sec. V, we will expand on avenues for further implementation of novel physics effects, and routes for extension to more general

HNL models. Finally, in Sec. VI, we will shortly discuss possible use cases for the generation app that could be of interest to the community, for HNL searches in present and future experiments.

II. THE HNL MODEL

In the framework of a seesaw-mechanism neutrino mass generation, an extended mass matrix \mathbb{M} sources the light neutrino masses [36,40,49]. Since the particle content of a seesaw mechanism generally includes new degrees of freedom in the neutrino sector (either Dirac states ν_R or Majorana states ν_L^c , with L, R indicating the chirality of the neutrino fields), the full lepton mixing matrix U has more than 3 rows:

$$U \in \mathbb{C}^{(3+n) \times (3+n)}, \quad \text{for } n \text{ new neutrino fields.} \quad (1)$$

These new fields are singlets under the Standard Model gauge group, which is to say, they are *sterile* neutrinos.

The most general Lagrangian admits both Dirac and Majorana neutrino mass terms [90],

$$\mathcal{L} \supset -\frac{1}{2} (\overline{\nu_L} \quad \overline{\nu_R^c}) \underbrace{\begin{pmatrix} 0 & m_D \\ m_D^T & m_M \end{pmatrix}}_{\mathbb{M}} \begin{pmatrix} \nu_L^c \\ \nu_R \end{pmatrix} + \text{H.c.} \quad (2)$$

Diagonalizing the full matrix in Eq. (2) yields $3+n$ eigenvalues corresponding to the 3 known light neutrinos, and n new neutrino fields. Heavy Neutral Leptons are mass eigenstates, with masses $m_N \gg 1 \text{ eV}/c^2$; they are therefore *nearly sterile*, and can interact with the Standard Model particle content through their mixing into the known neutrino flavor eigenstates $\alpha = e, \mu, \tau$,

$$\nu_\alpha = \sum_{i=1,2,3} U_{\alpha i} \nu_i + \sum_{j=4}^{3+n} U_{\alpha j} N_j. \quad (3)$$

A direct consequence of introducing new degrees of freedom is that the 3×3 PMNS mixing matrix $U^{\text{PMNS}} \subset U$ deviates from unitarity [91–93]. From solar and atmospheric neutrino oscillation analyses it is known that there are two nonzero mass splittings [94] $\Delta m_{21}^2 \simeq 7.4 \times 10^{-5} \text{ eV}^2$, $|\Delta m_{31}^2| \simeq 2.5 \times 10^{-3} \text{ eV}^2$. Though the sign of Δm_{31}^2 remains unknown, leading to two possible neutrino hierarchies, the fact remains that at least two mass eigenstates ν_i must have $m_{\nu_i} \neq 0$ (it is currently unknown whether the lightest mass eigenstate is massive or not). On the other hand, the most minimal HNL scenario, in which the lightest active neutrino is massless, corresponds to the case $n = 2$. In the most minimal extension to the Standard Model that admits neutrino masses (the νMSM , [38,42,49]), only Majorana mass terms are added; however, there exist other seesaw

mechanism implementations where Dirac masses are considered (see, e.g., [95,96]).

Following the approach of [53,57] we will work in the simplified framework of considering one extra neutrino state ($n = 1$), with the effective heavy neutrino N_4 controlled by the mixings $|U_{\alpha 4}|^2$. This approach is valid in the case of two HNLs with similar masses, assuming no coherent oscillations between them. (For a review of HNL oscillations, see, for example, [97,98] and references therein). The HNL in this minimal picture has no transition magnetic moment. There is, however, significant interest in the properties of neutrinos with transition magnetic moments, in the so-called “dipole portal” (the reader may refer to [99–101]). The general mixing for the HNL into flavor eigenstates Eq. (3) is now written as

$$\nu_\alpha = \sum_{i=1,2,3} U_{\alpha i} \nu_i + U_{\alpha 4} N_4. \quad (4)$$

In principle, the above information is sufficient to describe the production of N_4 from various different mechanisms. These include active-to-sterile neutrino oscillations [102], upscattering from neutrino-nucleus interactions in a material [103], decays of charged hadrons [58,104], or direct production at colliders [105–108].

We have implemented in GENIE the case of hadron decay into HNL, and the subsequent HNL decay into SM particles, within the framework of the effective field theory described in [104]. This theory describes the interactions of HNL with mesons directly at the diagram level, leading to HNL production and decay rates that are valid for HNL mass ranges up to about $1 \text{ GeV}/c^2$ and for energies of $\mathcal{O}(\text{GeV})$. Generally, it is possible to study the production and decays of HNL with larger masses, where production from heavy meson decays such as D, D_s as well as from τ and baryons is possible [109]. In this work, we have implemented HNL production from pions and kaons, as well as from muons, which are the dominant particles produced in current accelerator neutrino beamlines; we consider HNL with masses up to the kaon mass, for which there exist 10 lepton-number-conserving channels. We summarize the method used to obtain HNL production and decay rates in Appendix A 2.

At the end of this section, we provide for convenience the production and decay channels in tabular form. In Table I, we list the HNL production channels, along with the mass thresholds for channel activation, the “kinematic scaling” factor

$$\mathcal{K} = \frac{1}{|U_{\alpha 4}|^2} \frac{\Gamma(P \rightarrow N_4 + \ell_\alpha + \dots)}{\Gamma(P \rightarrow \nu + \ell_\alpha + \dots)}, \quad (5)$$

TABLE I. Production channels for HNL.

| Channel | Threshold (MeV/ c^2) | \mathcal{K} [110,111] | SM \mathfrak{B} |
|---|-------------------------|--|------------------------|
| $\pi^\pm \rightarrow N_4 + \mu^\pm$ | 33.91 | $\mathcal{P}_\ell(\delta_{m_\pi}^m, \delta_{m_{N_4}}^{M_{N_4}})$ | 0.999877 |
| $\pi^\pm \rightarrow N_4 + e^\pm$ | 139.06 | $\mathcal{P}_\ell(\delta_{m_\pi}^m, \delta_{m_{N_4}}^{M_{N_4}})$ | 1.23×10^{-4} |
| $K^\pm \rightarrow N_4 + \mu^\pm$ | 388.02 | $\mathcal{P}_\ell(\delta_{m_K}^m, \delta_{m_{N_4}}^{M_{N_4}})$ | 0.6352 |
| $K^\pm \rightarrow N_4 + e^\pm$ | 493.16 | $\mathcal{P}_\ell(\delta_{m_K}^m, \delta_{m_{N_4}}^{M_{N_4}})$ | 1.582×10^{-5} |
| $K^\pm \rightarrow N_4 + \mu^\pm + \pi^0$ | 253.04 | $\mathcal{S}_{K^+\mu}(M_{N_4})$ | 3.18×10^{-2} |
| $K^\pm \rightarrow N_4 + e^\pm + \pi^0$ | 358.18 | $\mathcal{S}_{K^+e}(M_{N_4})$ | 4.82×10^{-2} |
| $K_L^0 \rightarrow N_4 + \pi^\pm + \mu^\mp$ | 248.45 | $\mathcal{S}_{K^0\mu}(M_{N_4})$ | 0.2718 |
| $K_L^0 \rightarrow N_4 + \pi^\pm + e^\mp$ | 353.60 | $\mathcal{S}_{K^0e}(M_{N_4})$ | 0.3878 |

TABLE II. Decay channels for HNL. The pion decay constant $f_\pi = 130$ MeV.

| Channel | Threshold (MeV/ c^2) | Decay rate (Dirac) [104] | $\Gamma_{\text{Maj}}/\Gamma_{\text{Dirac}}$ |
|---------------------------------------|-------------------------|--|--|
| $N_4 \rightarrow \nu\nu\nu$ | 0 | $\frac{G_F^2 M_{N_4}^5}{192\pi^3} \sum_\alpha U_{\alpha 4} ^2$ | 2 |
| $N_4 \rightarrow \nu e^\pm e^\mp$ | 1.02 | $\frac{G_F^2 M_{N_4}^5}{192\pi^3} \cdot (\mathcal{C}_e(M_{N_4}) + \mathcal{D}_e(M_{N_4}))$ | 2 |
| $N_4 \rightarrow \nu e^\pm \mu^\mp$ | 106.17 | $\frac{G_F^2 M_{N_4}^5}{192\pi^3} \cdot U_{e4} ^2 \cdot \mathcal{F}(\delta_{M_{N_4}}^m)$ | $2 \frac{ U_{e4} ^2 + U_{\mu 4} ^2}{ U_{e4} ^2}$ |
| $N_4 \rightarrow \pi^0 \nu$ | 134.98 | $\frac{G_F^2 M_{N_4}^3 f_\pi^2}{32\pi} \sum_\alpha U_{\alpha 4} ^2 (1 - \delta_{M_{N_4}}^{m\pi^0})^2$ | 2 |
| $N_4 \rightarrow \pi^\pm e^\mp$ | 140.08 | $\frac{G_F^2 M_{N_4}^3 f_\pi^2 V_{ud}^{\text{CKM}} ^2}{16\pi} \mathcal{G}(\delta_{M_{N_4}}^m, \delta_{M_{N_4}}^{m\pi})$ | 2 |
| $N_4 \rightarrow \nu \mu^\pm \mu^\mp$ | 211.32 | $\frac{G_F^2 M_{N_4}^5}{192\pi^3} \cdot (\mathcal{C}_\mu(M_{N_4}) + \mathcal{D}_\mu(M_{N_4}))$ | $2 \frac{ U_{e4} ^2 + U_{\mu 4} ^2}{ U_{\mu 4} ^2}$ |
| $N_4 \rightarrow \pi^\pm \mu^\mp$ | 245.23 | $\frac{G_F^2 M_{N_4}^3 f_\pi^2 V_{ud}^{\text{CKM}} ^2}{16\pi} \mathcal{G}(\delta_{M_{N_4}}^m, \delta_{M_{N_4}}^{m\pi})$ | 2 |

and the Standard Model branching ratio \mathfrak{B} for each channel with $N_4 \rightarrow \nu$. By ν we mean the appropriate combination of $U_{\alpha i} \nu_i$ with α being the correct flavor for the channel. In Table II we list the decay channels, their mass thresholds, the expressions used to calculate the decay rate for a Dirac HNL (see Appendix A 2 for definitions of the expressions $\mathcal{C}, \mathcal{D}_\ell, \mathcal{F}, \mathcal{G}, \mathcal{P}_\ell, \mathcal{S}_{P\ell}, \delta_a^b$, and the multiplier used for the Majorana decay rates for each channel).

The production and decay “normalized branching ratios” of HNL are shown in Figs. 1 and 2. The normalization of these “branching ratios” is over all the available production (decay) channels with an HNL in the final (initial) state. This means that in Fig. 1, the sum of normalized branching ratios for each different parent is equal to 1, which is of course unphysical. The physical case of an HNL mixing weakly with the light neutrino sector implies that, to obtain HNL production rates, all these normalized branching ratios should be corrected by the overall factor [110]

$$\frac{\sum_{\alpha=e,\mu,\tau} |U_{\alpha 4}|^2}{1 - \sum_{\alpha=e,\mu,\tau} |U_{\alpha 4}|^2} \simeq \sum_{\alpha=e,\mu,\tau} |U_{\alpha 4}|^2, \quad (6)$$

where the denominator corresponds to the contribution from the active neutrino components.

Fundamentally, there is a difference between the behavior of two-body production channels (e.g., $K \rightarrow N_4 + \mu$, shown as “ $K \rightarrow \mu$ ” in the legend—the same convention is

applied throughout Fig. 1) and three-body channels (e.g., $K \rightarrow N_4 + \mu + \pi^0$). This is due to the breakdown of the helicity suppression mechanism that dominates the two-body production modes. In the Standard Model, the





FIG. 1. Normalized branching ratios for HNL production, $|U_{\tau 4}|^2 = 0$.


 FIG. 2. Branching ratios for HNL decays, $|U_{\tau 4}|^2 = 0$.

masslessness of the neutrino suppresses decays to electron + neutrino compared to muon + neutrino by a factor $\mathcal{O}(m_\mu^2/m_e^2)$ due to the neutrino being a helicity state as well as a left-chiral state. HNL, with masses $\gtrsim \mathcal{O}(100 \text{ MeV}/c^2)$, are not helicity eigenstates, and this suppression drops off with increasing mass, as shown in [110]. Three-body decays, on the other hand, have no helicity suppression mechanism, and the HNL production rate is instead controlled just by the reduction of phase space.

For the decay branching ratios, one generally has semileptonic two-body decays and (semi)leptonic three-body decays; however, the semileptonic three-body decays $N_4 \rightarrow \ell^\mp \pi^\pm \pi^0$, $N_4 \rightarrow \nu \pi^0 \pi^0$ have negligibly small branching ratios for $M_{N_4} < m_K$, confirming the argument made in [104] that these channels are essentially mediated by the emission of an on-shell ρ meson, $N_4 \rightarrow \ell^\mp \rho^\pm$, $\rho^\pm \rightarrow \pi^\pm \pi^0$ (and similarly for the two-neutral-pion decay). This assumption was checked using the *Mathematica* packages *FeynRules* [112], *FeynArts* [113], and *FeynCalc* [114–116] and the Lagrangian published on the *FeynRules* model database¹ [118]. The BR for the decay $N_4 \rightarrow \ell^\mp \pi^\pm \pi^0$ are $\lesssim 10^{-5}$, hence are irrelevant for the mass range $M_{N_4} < m_K$.

III. PREVIOUS IMPLEMENTATIONS OF HNL SIMULATION

To fully exploit the wealth of parameter space that is accessible to accelerator neutrino experiments in searching for HNL, it is necessary to develop some commonly accepted prescription to produce well-understood and validated simulations of HNL. Of special interest is the

¹We used the relevant Feynman diagrams (see Fig. 16 in Appendix A 2) to get the matrix element for the $N_4 \rightarrow \ell^\pm \pi^\pm \pi^0$ decay, and obtained a differential decay rate $d^2\Gamma/dE_\ell dE_{\pi^\pm}$, which is integrated using Simpson’s rule on the nested integral (see [117] Lecture 24).

ability to adequately describe any phenomena that arise from the massive nature of the HNL, particularly in the kinematics, as this could lead to powerful methods to remove background contributions, such as with a suitable timing trigger (see, e.g., [54,71]) or using appropriate final-state observables [53,57]. The suite of neutrino event generators commonly used by neutrino experiments to model neutrino-nuclear interactions [83,119–121] have in common the desirable trait of each having a unified interface for all the various physics of Standard Model neutrinos [80]. This allows for systematic and robust study of the complex nuclear environment that governs these interactions.

For BSM physics, though there has been considerable activity (particularly in the context of colliders) in unifying the production pipeline on the theory and phenomenology side [67,104,122,123], progress for the few-GeV range of neutrinos is less rapid. Codes to handle the production of HNL from a neutrino beam and their decays to visible final states, with varying degrees of complexity, have been implemented using either existing software frameworks such as *PYTHIA* [72,124] and *GEANT4* [104], purpose-built new simulation tools [54,70,82,125], or applying an appropriate weighting procedure to existing beam simulations [57].

There is an increasing need, then, to aid analyses aiming to look for HNL on two fronts: on the one hand, different analyses should be able to be compared straightforwardly for reproducibility and clarity; on the other, analysis resources should be conserved as much as possible, focusing on leveraging the unique capabilities of next-generation detectors for rare searches. The current landscape of individual implementations of HNL in each experiment certainly stands to gain from such capacity, not least in the case of interfacing the work of theorists, phenomenologists, and experimentalists who are working toward the same goal.

A natural expansion into the requirements outlined above is to adapt this new physics into commonly used generators,

and in the case of HNL produced by neutrino upscattering has already been initiated both in GENIE [88] and in the recent DarkNews generator [82]. For the case of HNL from hadron decays, the event generation philosophy is different: instead of producing neutrino interactions based on the *same* flux as for Standard Model neutrinos, HNL decays are produced from an *a priori* very different flux, which results from the kinematics and dynamics of the massive neutrino; apart from the trivial scaling caused by the mixings to the SM, there is a dynamic shape distortion sourced by the generally slower, unstable particles that must reach the detector and decay inside without traversing it. Furthermore, because HNL in this paradigm are produced directly from the hadron spectrum engendered in proton-target interactions, it is not sufficient merely to rescale the amount of recorded protons-on-target (POT) to estimate HNL event rates, or to reweight the SM neutrino flux; a full description of each HNL must be capable of tracking how many POT are being considered to avoid normalization errors in the analysis downstream.

IV. IMPLEMENTATION IN GENIE v3

We extend the effort being built by contributing a detailed simulation of HNL produced by hadron decays inside a neutrino beamline within the GENIE framework, and their subsequent decays in a detector, with a view particularly towards present and future HNL search efforts [31,69,71,77,126].

We start by describing three coordinate systems, *de facto* used in an accelerator neutrino experiment. A sketch of a typical neutrino beamline outlining these coordinate systems, and showing how an HNL is typically produced, is shown in Fig. 3.

First, a “NEAR” frame (x_N, y_N, z_N) defines the beamline coordinates. This is typically set such that the target hall building has its floor parallel to $y = \text{const}$ and such that the origin O is within the target. Second, a “BEAM” frame (x_B, y_B, z_B) , normally obtained by rotating (x_N, y_N, z_N) downwards in the (y_N, z_N) plane.

The z_B axis is typically parallel to the direction of the neutrino beam. Third, a “USER” frame (x_U, y_U, z_U) , attached to a detector some distance away from the target. For example, take the case of DUNE [127,128], assuming the liquid argon component to be in the on-axis configuration and at a distance of 575 m from the origin O, and that the LBNF beam is rotated 5.8° downwards, the positions of




FIG. 3. Cartoon of a typical neutrino beamline. See text for details.

the detector center in NEAR, BEAM, and USER coordinates are, respectively,

$$\begin{cases} (x_N, y_N, z_N) = (0, -58.11, 572.06) \text{ m}, \\ (x_B, y_B, z_B) = (0, 0, 575) \text{ m}, \\ (x_U, y_U, z_U) = (0, 0, 0). \end{cases} \quad (7)$$

Similarly, for the MINERνA inner detector (tracker + nuclear targets) [129], the NEAR, BEAM, and USER positions are taken to be

$$\begin{cases} (x_N, y_N, z_N) = (-0.25, -60.35, 1022.74) \text{ m}, \\ (x_B, y_B, z_B) = (-0.25, -0.66, 1024.52) \text{ m}, \\ (x_U, y_U, z_U) = (0, 0, 6.44) \text{ m}. \end{cases} \quad (8)$$

For a detailed explanation of the convention used for positioning, see Appendix A 3.

It is instructive at this point to write down the ideal factorization for such an endeavor:

- (1) A dynamic HNL flux prediction, taking into account the detailed kinematic effects of particles traveling with $\beta < 1$.
- (2) A sophisticated decay library, including effects of HNL polarization.
- (3) A complete array of bookkeeping tools, including POT counting.

A preeminent position must be given to the quite general flux and geometry drivers implemented in GENIE, which allow this implementation to take advantage of the full complexity of a hadron beam simulation and geometry description, making this tool suited to studies at various stages of maturity of an experiment from its very early stages all the way through to its most mature stages. Conceptually, the GENIE BeamHNL module may consume an arbitrarily complex simulation of hadrons in some beamline as input, along with an arbitrarily complex detector description, and output HNL decay events within the desired detector volume, in user-defined coordinates. This underlines the nature of this tool as a self-contained kit to handle HNL kinematics and dynamics from production up until decay, with no assumptions *a priori* about the nature of the experiment it is being applied to.

A. HNL production in beamline

The starting off point for deriving a neutrino flux is a precalculated spectrum of particles, such as is typically produced from proton interactions on a target in a neutrino beamline. These particles (such as π^\pm, K^\pm, μ^\pm) are then usually focused by a series of magnets to perform some charge selection. The particles then propagate downstream, decaying in some suitably long decay volume to neutrinos (or particles that eventually decay to neutrinos). They are referred to as “parents” if they decay directly to neutrinos, or

more generally “ancestors” if one of their decay products is a parent. One or more particle absorbers are normally placed in between the decay volume and the detector, so that the only particles that survive to reach the detector are neutrinos. This same production philosophy applies to HNL, which by Eq. (4) have a probability $\propto |U_{\alpha 4}|^2$ to be the mass state that corresponds to the ν_α flavor state made during neutrino production.

Experiments simulate the spectrum for each parent species using a description of their beamlines and suitable external experimental data ([130,131]) to constrain the simulation uncertainty [132]. There are various formats the output of this simulation can be stored in; we have chosen to adapt the `dk2nu` format [133] developed for the Fermilab Intensity Frontier experiments to a “`flatdk2nu`” format that mirrors the `dk2nu` tree structure, without containing any complex classes. This was done to minimize the build complexity for GENIE. However, other flux formats can readily be converted into the format required for this simulation, and an example input flux with the necessary structure has been provided in the `$GENIE/src/contrib/beamhnl` directory accompanying our module.

One application of this general format is the ability to use this module not only for accelerator neutrino experiments such as [28,69,126], but also in the context of higher-energy collider neutrino experiments with detectors lying downstream of the neutrino production point [72,77]. This is particularly interesting, because at higher energies heavy mesons (such as D, D_s), which can decay to HNL heavier than the kaon mass, are produced copiously enough to have a strong sensitivity to HNL with $M_{N_4} > m_K$. Equally, the same format could be adapted for use with atmospheric neutrino fluxes [58,134] to probe HNL of extraterrestrial origin. We comment further on this in Sec. V.

We have used the production branching ratios from [110,111] for the evaluation of HNL production probabilities, for the most commonly produced parent types: π^\pm, K^\pm, K_L^0 (directly from the beamline), and μ^\pm (from hadron decay themselves). Further hadron types are not yet a part of our simulation, but are compelling targets to include.

To conserve computing resources, we force every parent in the simulation to decay to HNL; that is equivalent to using the uncorrected branching ratios from Fig. 1. This is corrected in the simulation bookkeeping by applying Eq. (6) on the production probability of the HNL.

A salient feature of the HNL production is that, depending on the mass of the HNL, the acceptance of the detector as seen by the HNL at its production point can change significantly, owing to the fact that HNL travel with $\beta_{N_4} < 1$. A more detailed review of massive-particle kinematics can be found in [135].

Consider a parent P of energy E_P , decaying leptonically to a neutral lepton L that can either be light (ν) or heavy

(N_4) and a charged lepton ℓ . The energy of L in the rest frame of P is then

$$E_L^* = \frac{m_P^2 - m_\ell^2 + m_L^2}{2m_P}, \quad (9)$$

Suppose, without loss of generality, that an observer sees P propagate along the z axis, $\mathbf{p}_P = (0, 0, p_P)$. Then the angle Θ at which L is emitted with respect to the z' axis in the rest frame can be related to the emission angle θ in the observer’s frame, as

$$\tan \theta = \frac{q_L \sin \Theta}{\gamma_P (\beta_P E_L^* + q_L \cos \Theta)}, \quad (10)$$

where q_L is the rest-frame momentum of L , and β, γ are the relativistic parameters of P in the lab frame. Figure 4 shows θ as a function of Θ for a kaon parent, with $E_K = 1$ GeV, for various HNL masses produced in the decay $K^\pm \rightarrow N_4 + \mu^\pm$.

The cases of massless neutrinos and HNL are quite different; for massless Standard Model neutrinos, $q_L = E_L^*$ and the resulting function $\tan^{-1} [\gamma_P^{-1} \sin \Theta (\beta_P + \cos \Theta)^{-1}]$ is monotonically increasing. This means that, for any arbitrary emission angle θ_0 in the lab frame, there exists some suitable preimage Θ_0 in the rest frame, for any parent velocity. In a Standard Model calculation, *any* parent can produce a neutrino that is accepted by a detector, *for any* parent momentum and relative position of the neutrino production vertex to the detector.

In contrast, for HNL with masses large enough, it is not generically true that any θ_0 can be achieved; this places increased importance on the kinematics of the parent P . The greater the angle between P ’s momentum \mathbf{p}_P and the relative separation \mathcal{O} between neutrino production vertex and




FIG. 4. Collimation effect for HNL in $K^\pm \rightarrow N_4 \mu^\pm$, $E_{K^\pm} = 1$ GeV. For M_{N_4} heavy enough, there is no backward emission in the lab frame.

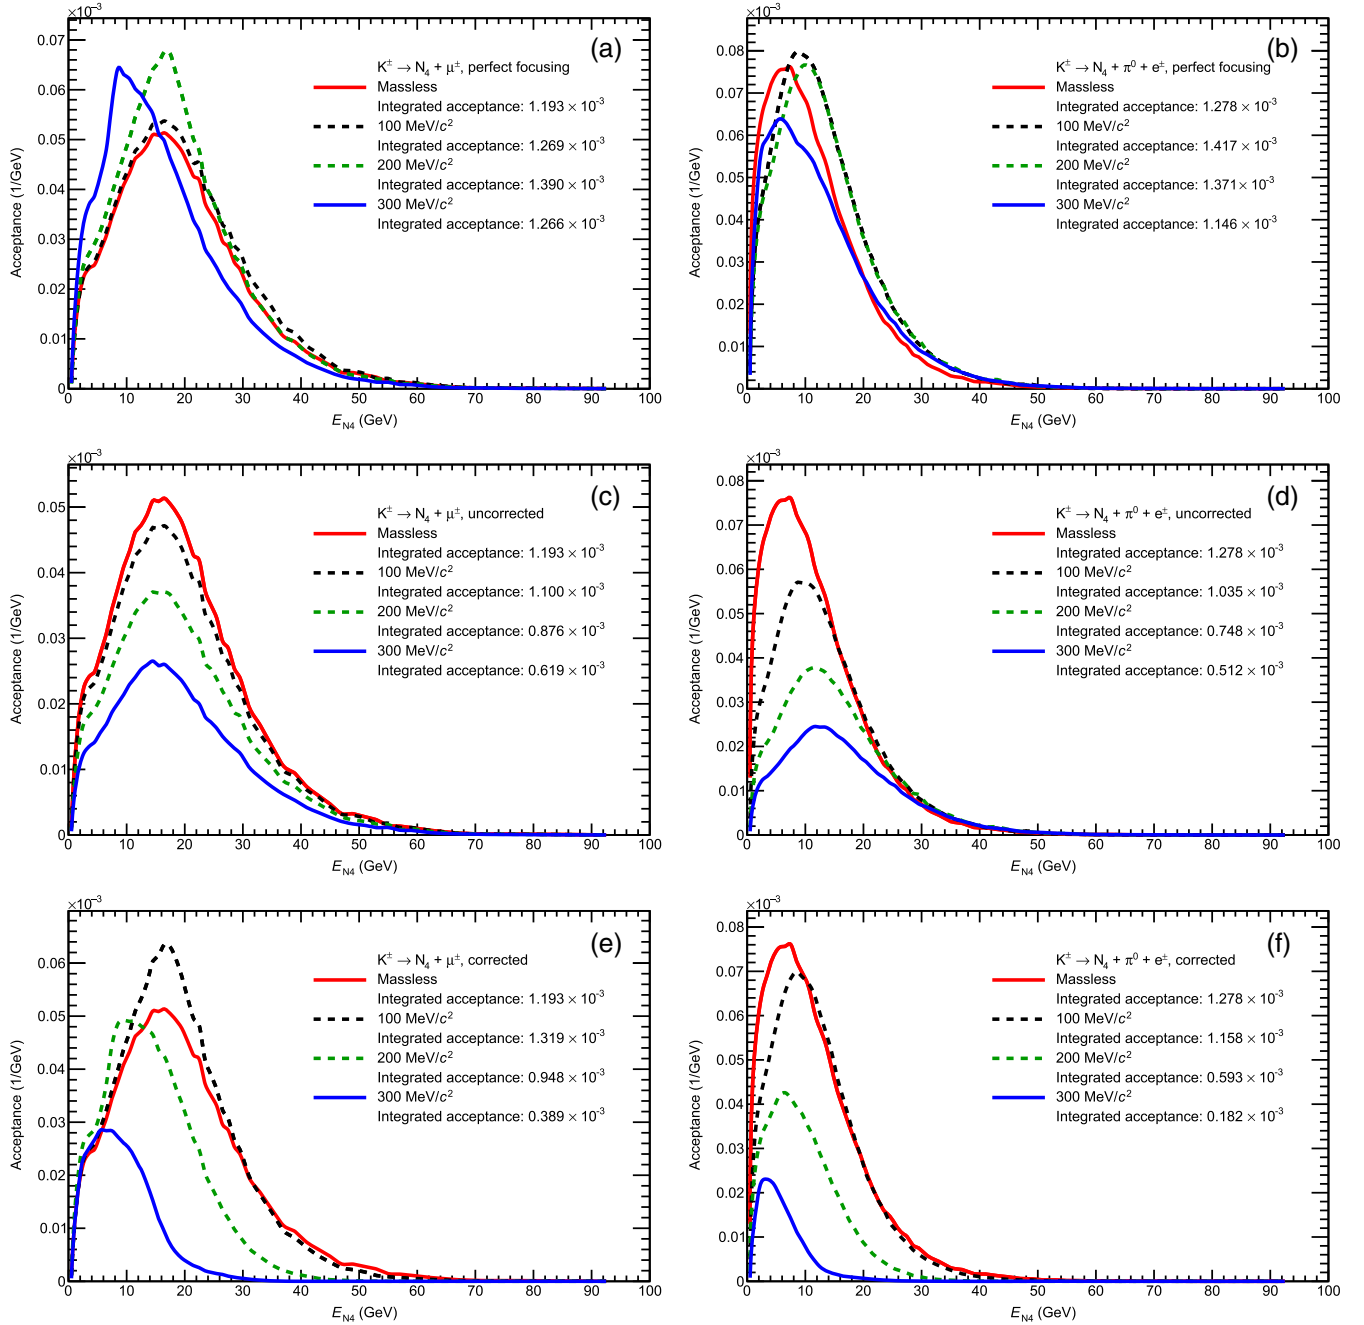


FIG. 5. Acceptance for a K^\pm decaying to HNL and either μ^\pm or $e^\pm + \pi^0$, for the MINER ν A detector [129]. Top row: Parents perfectly focused: momentum of parent always points towards detector. Middle row: parents not perfectly focused, but no acceptance correction \mathcal{A} applied. Bottom row: full focusing and acceptance correction applied. The normalization of these curves drops considerably when focusing is not perfect; this happens because the boost factor \mathcal{B} drops with opening angle between parent momentum and detector location. The effect of \mathcal{A} is mainly to suppress the high- E_{N4} tails, as hard HNL cannot deviate from their parents enough to reach the detector; also, at low enough HNL mass, \mathcal{A} increases normalization as backward-emitted HNL are accepted by the detector.

detector, and the greater the HNL mass, the smaller the acceptance becomes, generally speaking; for large enough angles and masses, the HNL could not be accepted at all, regardless of its rest-frame emission angle. This same collimation effect can, at the same time, account for *increased* acceptance of HNL with respect to a Standard

Model neutrino, if P is well-collimated enough. This happens because Eq. (10) now reaches a maximum and tends towards 0 for large Θ : in other words, backwards-emitted HNL are swept forwards by the Lorentz boost into the lab frame, and end up accepted by the detector. In general, this *acceptance correction* respective to a Standard

Model neutrino is calculated on an event-by-event basis, and applied as an additional weight to each HNL. One can also obtain an estimate for the neutrino energy at the detector, again setting $p_L^x = 0$ without loss of generality

$$\begin{aligned} \Rightarrow E_L^* &= \gamma_P E_L - \gamma_P \beta_P p_L \cos \theta_D, \\ \Rightarrow E_L &= \frac{E_L^*}{\gamma_P (1 - \beta_P \cos \theta_D)} \equiv \mathcal{B} E_L^*, \end{aligned} \quad (11)$$

where \mathcal{B} is termed the *boost factor*, and θ_D is the viewing angle between \mathbf{p}_P and \mathcal{O} . From a Standard Model simulation standpoint, this equation simplifies to

$$\mathcal{B}_\nu = \frac{1}{\gamma_P (1 - \beta_P \cos \theta_D)}, \quad (12)$$

which uniquely determines the lab-frame energy for the massless neutrino. For a massive neutrino, there is a complication: Eq. (11) depends on knowledge of the lab-frame velocity through β_L . We estimate β_L by imposing a geometric constraint; using the worldline (T, \mathcal{O}) with $T = |\mathcal{O}|/(\beta_L c)$, we construct a candidate lab-frame momentum by boosting (T, \mathcal{O}) into P 's rest frame and forcing the HNL momentum to point to that direction, then boosting the result back into the lab frame. We check the distance between the point of closest approach and the detector center; if this is too large, we decrement β_L and repeat the procedure.

In Fig. 5, we plot the differential geometrical acceptance, defined as the probability a HNL emitted isotropically in the parent's rest frame will be accepted by the detector, for the case of the MINER ν A detector [129] in the NuMI medium-energy beam [132]. We have plotted the acceptance for the processes $K^\pm \rightarrow N_4 + \mu^\pm$ (left-hand side plots) and $K^\pm \rightarrow N_4 + \pi^0 + e^\pm$ (right-hand side plots). Panels (a), (b) show the differential acceptance under the assumption the parent kaons are perfectly focused, which is to say an HNL emitted with momentum collinear to the kaon momentum would definitely be accepted. Generally, the acceptance remains about the same as for Standard Model neutrinos, though the peak shifts first to higher and then to lower energies. This is caused by the Lorentz boost becoming more efficient for all HNL as the mass initially increases, followed by the decrease associated with the drop in boost factor \mathcal{B} as the HNL's velocity drops significantly. In panels (c), (d), we have used the kaon spectrum from the NuMI beamline simulation including realistic focusing, but not applied acceptance correction. The shape of the acceptance remains roughly similar, but the normalization and integrated acceptance change dramatically. This is caused by the suppression of the boost factor \mathcal{B} at both high angles and low velocities. Finally, panels (e), (f) show the full simulation accounting for both realistic focusing and the change in suitable emission regions in the parent rest frame due to acceptance correction. Peaks of these distributions shift decidedly to




FIG. 6. Effect of focusing on acceptance; flux at MINER ν A (realistic focusing)/flux (perfect focusing). Error bands are purely statistical.

lower energy, as the higher energy HNL are more collimated with their parents and, unlike Standard Model neutrinos, are not necessarily able to reach the detector. This effect becomes increasingly prominent as M_{N4} goes up and the kinematic constraints become more severe.

We further demonstrate the effect of parent focusing on the acceptance in Fig. 6, by simulating the flux under the assumption of perfect parent focusing or using realistic parent focusing as provided from the NuMI beamline simulation. We show the ratios of realistic over perfect acceptance, as a function of HNL energy, for masses 0, 100, ..., 400 MeV/ c^2 . The effect is most pronounced at 100 and 400 MeV/ c^2 , where the thresholds of HNL production by pions and kaons are almost reached; at these masses, the collimation effect becomes most severe, and proportionally more parents can not produce HNL that are accepted by the detector, unless the parent happens to be traveling in a direction that would intersect the detector. The effect is more pronounced at high HNL energy E_{N4} , because more energetic HNL have a stronger collimation effect due to the larger Lorentz boost. Under realistic focusing, the HNL spectra become softer; this is due to the suppression of the boost factor with parent momentum angle.

Another profound impact can be seen on off-axis neutrino spectra, relevant for PRISM-like searches. Varying Eq. (11) over $\cos \theta$ for $\beta_L = 1$ yields the celebrated *off-axis effect* for the Standard Model. In the general case $\beta_L \neq 1$, however, the off-axis effect becomes progressively less pronounced. In Fig. 7, we have chosen a pion parent, and four different values of β_L ; the Standard

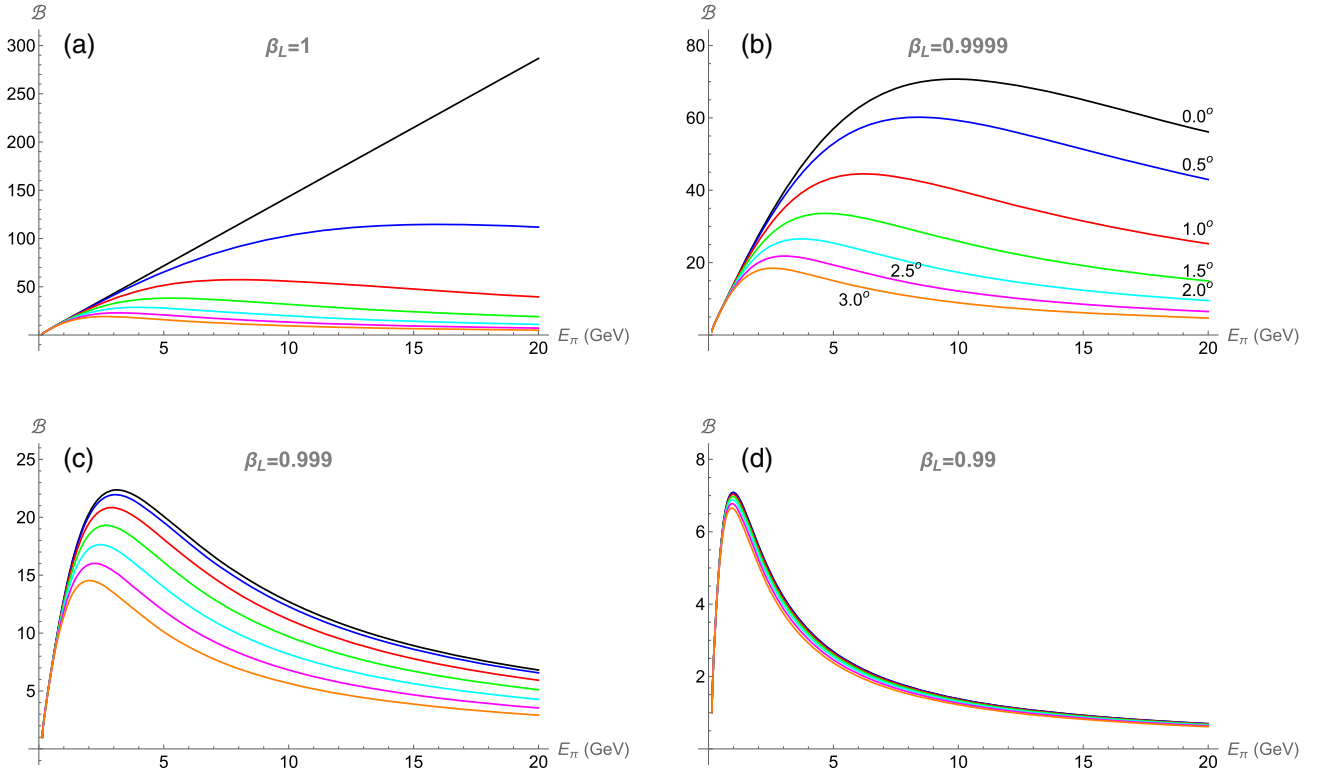


FIG. 7. Boost factor \mathcal{B} as function of pion energy, for different values of θ_{OA} . (a) $\beta_L = 1$ (Standard Model); (b) $\beta_L = 0.9999$; (c) $\beta_L = 0.999$; (d) $\beta_L = 0.99$. The off-axis angle θ_{OA} grows from top to bottom, from 0.0° to 3.0° in increments of 0.5° .

Model case is retrieved in (a), and progressively smaller β_L are shown in (b–d). The various curves representing different values of θ_D collapse to a single curve. Writing out the derivative $\partial E_L / \partial E_P = E_L^* \partial \mathcal{B} / \partial E_P$, one can find the peak \mathcal{B} to be at

$$E_P|_{\mathcal{B}_{\text{max}}} = \frac{m_P}{(1 - \beta_L^2 \cos^2 \theta)^{1/2}}. \quad (13)$$

In Fig. 8 we have plotted, for different values of θ_D and for pion parent, $E_P|_{\mathcal{B}_{\text{max}}}$ (x axis) as a parametric plot of $\log_{10}(1 - \beta_L)$ (y axis). Standard Model neutrinos are the limit $y \rightarrow \infty$; additionally, we have drawn two axes for visualization: red (left) is the mass of a 1 GeV HNL for the given β , and blue (right) is the energy a $100 \text{ MeV}/c^2$ HNL would have at that β . For very high β , the position of the peaks varies greatly with E_π , meaning the off-axis effect is visible; this motivates, for example, the PRISM concept [128] of utilizing a narrower neutrino spectrum at high off-axis angles to constrain the neutrino flux. However, as β decreases, so does the variation of the neutrino spectrum [as one can also see in Fig. 7 panel (d), with curves converging into one for $\beta = 0.99$]: since β scales inversely to M_{N_4} , the HNL flux is expected to be *less* sensitive to the off-axis effect than Standard Model neutrinos. In principle, the spectrum of HNL is still

expected to become softer [70], but the effect seen would be much smaller than predicted previously. In Fig. 9, one can see our calculation of the flux shapes² at the DUNE PRISM with $|U_{e4}|^2 : |U_{\mu 4}|^2 : |U_{\tau 4}|^2 = 1 : 1 : 0$, for masses $M_{N_4} = 0, 100, \dots, 400 \text{ MeV}/c^2$. Notice how as the off-axis angle θ_{OA} changes, the massless neutrino flux shifts in accordance with the SM prediction, whereas the $M_{N_4} = 100, \dots, 400 \text{ MeV}/c^2$ HNL have a far smaller (but still present) dependence on the off-axis angle.

We have further shown in Fig. 10 an example of the diminished off-axis effect on the truth-level muon energy distribution from the decay $N_4 \rightarrow \pi^+ + \mu^-$ with HNL mass $M_{N_4} = 300 \text{ MeV}/c^2$ and $|U_{e4}|^2 = |U_{\mu 4}|^2$, and compared with the expectation on the equivalent distribution of charged-current muon neutrino interactions with one μ^- in the final state. This would correspond to panel (d) of Fig. 9, where the flux of HNL that are crossing the detector is shown, for the bins $[0, 1)$, $[15, 16)$, and $[30, 31)$ in OA displacement. In the first three panels of Fig. 10, the suppression of the off-axis effect on HNL energy distributions (black curves) at the detector, which is taken to be a 5 m-side cube at the coordinates (7), is apparent, especially when compared to the Standard Model muon

²We have used the DUNE flux files available at [136], for the ‘‘Optimized 3-Horn Design with 1.5 m target and Fully Engineered Horn A (January 2021)’’ configuration.




FIG. 8. Pion energy at which \mathcal{B} is highest, for different values of θ .

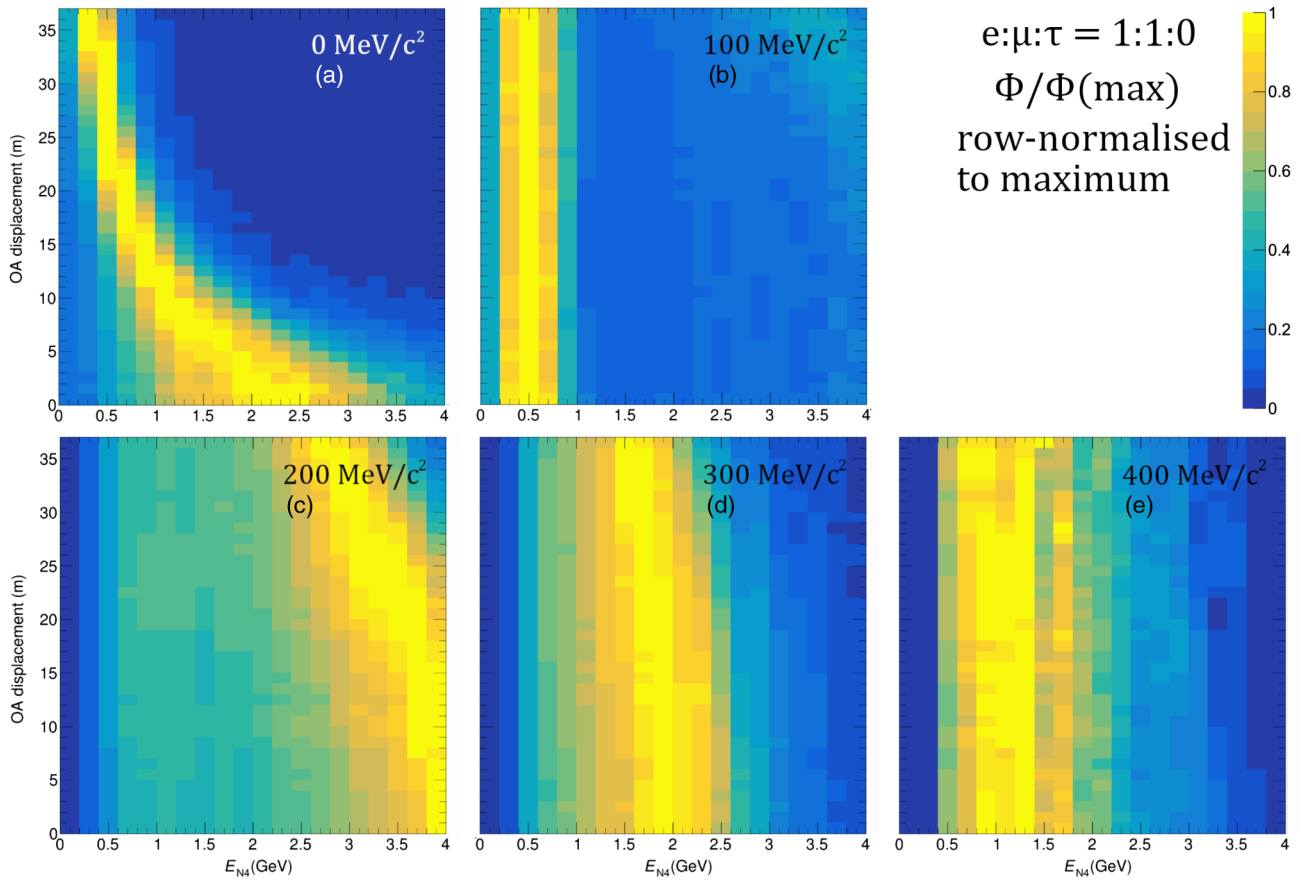


FIG. 9. HNL flux shapes at DUNE PRISM, as function of HNL energy and off-axis displacement.

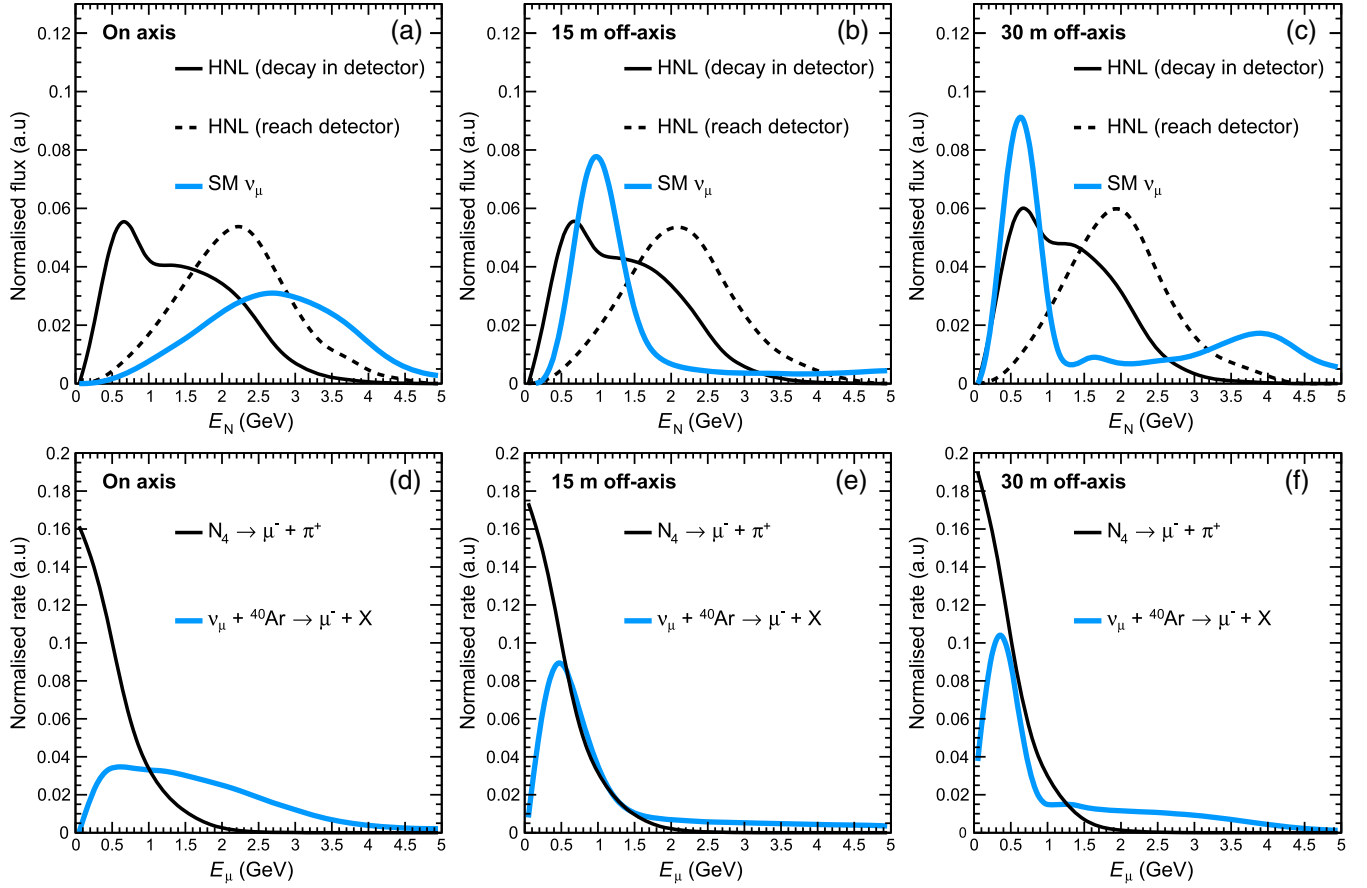


FIG. 10. Example of truth-level muon spectral shapes originating from the HNL decay $N_4 \rightarrow \pi^+ + \mu^-$ ($M_{N_4} = 300 \text{ MeV}/c^2$) in a 5 m-side box-shaped detector at the DUNE PRISM baseline, at three off-axis displacements. Panels (a–c) show the expected flux shape of HNL that decay inside the detector (solid black line) and that reach the detector (dashed black line), compared to the expected flux of Standard Model muon neutrinos (solid blue line). Panels (d–f) show the muon energy from HNL decays (black) vs the expected muon spectrum from charged-current inclusive interactions of ν_μ on ^{40}Ar . SM interactions simulated with G18_02a_00_000.

neutrino expectation (solid blue curves). Note that the dashed curves, which show the HNL that are expected to cross the detector but not necessarily decay inside it, are markedly different to the solid curves which correspond to HNL actually decaying inside the detector. This is due to the lower velocities of softer HNL causing them to decay in the detector at a disproportionately high rate. Also note the difference from the solid blue curves that represent the SM expectation, in line with the suppression of the off-axis effect for slower-than-light particles as shown in Figs. 7 and 8. For the Standard Model muon energy distributions we have used events generated with GENIE version 3.02.00 on argon-40, selecting only those events with precisely 1 muon in the final state.

The baseline for the detector is 575 m, which means the reference off axis transverse displacements of 5, 10, 20, and 30 m are equivalent to $\theta_{\text{OA}} \simeq 0.5, 1.0, 2.0, 3.0^\circ$.

Our HNL simulation produces the flux calculation based on minimal information from the beamline simulation, which is passed as an input.

There are two fundamentally important inputs:

- (1) Parent momentum and decay vertex position in NEAR coordinates.
- (2) “Importance weight” [56,133]—a multiplicity factor for hadrons with very similar kinematics.

Based on this information, the module assigns the appropriate decay channel, calculates the boost factor to obtain the energy of the HNL at the detector under the constraint that the HNL can reach it, and constructs the first particle in GENIE’s particle stack that corresponds to an HNL that decays in the detector. We provide details of the book-keeping in Appendix A 4.

B. Decay to Standard Model particles

Unlike Standard Model neutrinos, HNL are unstable and can decay directly (semi-)leptonically to SM particles.

The lifetime of an HNL is generally inversely proportional to the mixing with the SM leptons $\sum_\alpha |U_{\alpha 4}|^2$ and to the HNL mass M . Depending on the parameter space point being searched, the branching ratios for HNL decays vary as

thresholds for various channels open; a few possibilities can be seen in Fig. 2. Details of how we keep track of the decay channels and selects the correct one are in Appendix A 2. For HNL below the kaon mass, the most prevalent production channels are the two-body production channels $\pi \rightarrow N_4 + \ell$ and, above the pion threshold, $K \rightarrow N_4 + \ell$; for HNL above about $100 \text{ MeV}/c^2$, the main decay channels are the two-body channels $N_4 \rightarrow \pi + \ell, N_4 \rightarrow \pi^0 + \nu$.

Since the SM weak force couples to left-chiral particles and right-chiral antiparticles, the decays of N_4 and \bar{N}_4 have opposite angular dependencies [97]; the well-known corollary is that in two-body decays only Dirac HNL have a $\cos\theta$ dependence in their decay spectra, whereas Majorana HNL do not (as has been shown explicitly in the case of neutral-mediated decays in [137]). In practical terms, charge-blind detectors that cannot distinguish between leptons and antileptons cannot search for forward-backward asymmetries in the decay distributions of HNL [53].

An HNL has the same intrinsic angular momentum as a Standard Model neutrino; however, its mass implies that it is *partially*, rather than completely, polarized [138]. Equivalently, the HNL and the other particle(s) that were produced in the decay of the pseudoscalar P form a pure $J=0$ state, but the HNL itself is not a pure state. Considering the leptonic production mode $P \rightarrow N_4 + \ell'$, one can write down a polarization vector \mathbb{P} for the HNL. \mathbb{P} has magnitude $|\mathbb{P}| < 1$ (partial polarization) and direction collinear to the momentum $\mathbf{q}_{\ell'}$ where \mathbf{q} is written in the HNL's rest frame. Because in P 's rest frame the momenta \mathbf{p}_{N_4} and $\mathbf{p}_{\ell'}$ are collinear, \mathbb{P} 's direction is the same as the momentum \mathbf{q}_P of P in the HNL rest frame.

It is important, moreover, to keep track of where the polarization vector \mathbb{P} of the HNL is pointing, since it defines the only ‘‘privileged direction’’ in the HNL rest frame, and thus the axis with which an angular dependence of the decay products $N_4 \rightarrow X$ manifests. \mathbb{P} 's direction can be tracked at the moment of HNL production.

The partial polarization of N_4 also means that its differential decay rate into the generic final state X is multiplied, again, by the appropriate coefficient determined by $|\mathbb{P}|$. Essentially, instead of the well-known $1 \mp \cos\theta$ dependence expected from the decay of fully polarized spin-1/2 particles, the angular dependence is modulated by a suitable polarization modulus H that is *a priori* dependent on both the production and decay modes of the HNL, resulting in an angular dependence $1 \pm H \cos\theta$. The reader will notice that $-1 < H < 1$, i.e., the polarization modulus can switch sign depending on the HNL mass. Its sign, s_{N_4} , is $+1(-1)$ if the expectation value of measuring the HNL spin on the direction of its momentum is positive (negative). In the simple case of two-body production and two-body decay $P \rightarrow N_4 + \ell', N_4 \rightarrow D + \ell$, H is explicitly given as [138]

$$H = - \frac{(m_{\ell'}^2 - M_{N_4}^2)\lambda^{1/2}(m_P^2, M_{N_4}^2, m_{\ell'}^2)}{m_P^2(M_{N_4}^2 + m_{\ell'}^2) - (m_{\ell'}^2 - M_{N_4}^2)^2} \cdot \frac{(M_{N_4}^2 - m_{\ell}^2)\lambda^{1/2}(M_{N_4}^2, m_{\ell'}^2, m_D^2)}{(M_{N_4}^2 - m_{\ell}^2)^2 - m_D^2(M_{N_4}^2 + m_{\ell}^2)} = (s_{N_4} \cdot |\mathbb{P}|) \cdot |\mathbb{D}|, \quad (14)$$

where $\lambda(x, y, z)$ is the Källén function defined in Appendix A 2, $s_{N_4} = +1(-1)$ if M_{N_4} is greater (smaller) than $m_{\ell'}$, and \mathbb{D} is the factor resulting from the decay $N_4 \rightarrow D + \ell$. The factorization of the polarization modulus from HNL production and decay is thus made apparent.

In this work, we have implemented the simple, yet analytically calculable ‘‘two-body-production, two-body-decay’’ polarization prescription for all (Dirac) HNL decays, assigning an angular distribution $1 \pm H \cos\theta$ to the spectrum of decay products. However, we have implemented a switch that allows the user to turn this simple scheme off if it is desirable to do so, reverting to pure phase-space decays instead. Because the user has access to the truth-level four-momentum of each of the decay particles, it is in principle possible to reweight the spectra of the final state products with any desired polarization scheme. The implementation of a fuller description of polarization effects, including in the three-body decays of Majorana HNL (see, for example, [111, 139] for discussions on Majorana HNL polarizations), remains an appealing avenue for future work. We comment on this in Sec. V.

We have simulated in Fig. 11 the effect of three different polarization prescriptions for $400 \text{ MeV}/c^2$ HNL from the kaon decay $K^+ \rightarrow N_4 + e^+$, decaying as $N_4 \rightarrow \pi^+ + \mu^-$ at the MINER ν A detector in the NuMI beam [129]. The modulus H is plotted as a function of M_{N_4} in Fig. 12, with the red curve corresponding to the production mode $K^+ \rightarrow N_4 + e^+$ and the blue curve to $K^+ \rightarrow N_4 + \mu^+$; for $M_{N_4} = 400 \text{ MeV}/c^2$ and $K^+ \rightarrow N_4 + e^+$, it is $H \simeq 0.9961$. The simplest prescription in Fig. 11 (blue curve) is the absence of any polarization effect; the angular distribution of final-state products is isotropic in the HNL rest frame, as

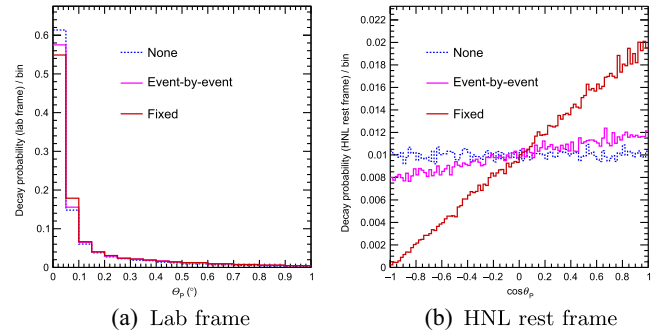


FIG. 11. Angular distributions of final-state muons in $K^+ \rightarrow N_4 + e^+, N_4 \rightarrow \pi^+ + \mu^-$ for Dirac HNL with $M_{N_4} = 400 \text{ MeV}/c^2$, using NuMI flux and for a detector located at MINER ν A's coordinates. See text for details.




FIG. 12. Polarization modulus H for the chains $K^\pm \rightarrow N_4 + e^\pm, N_4 \rightarrow \mu^\mp + \pi^\pm$ (red), $K^\pm \rightarrow N_4 + \mu^\pm, N_4 \rightarrow \mu^\mp + \pi^\pm$ (blue).

expected. In the red curve, a “maximal scenario” of polarization is implemented. The *direction* of \mathbb{P} is kept fixed to \hat{z} , which shows up as a distribution $1 + H \cos \theta_{\mathbb{P}}$, where $\theta_{\mathbb{P}}$ is the angle between the muon and \hat{z} . This shows that, for suitable HNL mass, there could be a significant polarization effect if one did not account for the variation of the direction of \mathbb{P} . Finally, the magenta curve implements the realistic scenario where the direction of \mathbb{P} is evaluated event by event according to the procedure outlined above Eq. (14). It is immediately apparent that the polarization is almost completely washed out; this occurs as a consequence of backwards-emitted HNL being accepted by the detector. For such a heavy mass, there is almost the same chance of a backwards-emitted HNL being accepted as a forwards-emitted one; this “averages out” the polarization effect.

The proper implementation of polarization effects is doubtlessly going to be crucially important for lower-energy beamlines, or decays at rest, where the Lorentz transformations into the lab frame are small or identity; for the case of the 120 GeV NuMI beam, for example, the transformation into the lab frame makes polarization effects matter very little for the correct description of final-state kinematics, affecting the opening angle $\Theta_{\mathbb{P}}$ between the muon momentum and beam direction in the region $\Theta_{\mathbb{P}} \lesssim 0.1^\circ$.

C. Determination of decay vertex

The final task for the description of a single HNL decay event is its location in spacetime, given some origin; frequently, the user will input their own coordinate system, which is just the USER frame defined earlier.

Given the velocity β and lifetime τ of the HNL, one can calculate where in the detector the decay occurs, which is of

unique interest to segmented detectors such as the DUNE near detector [128], whose different submodules may have different tracking capability, thresholds, geometries and fiducial volumes, etc.

Since fast HNL also spend less time in the detector, there is an effect of the HNL velocity (and hence parent velocity) on the expected spectrum of decays; in other words, the detector dimensions and location are convolved with the flux information [140].

Heavier HNL tend to “lag behind” Standard Model neutrinos, and thus could, for long enough baselines and small enough velocities, decay during a timing window with little to no Standard Model expected background such as the period in between beam spills, or even in between beam bunches ([141,142]). This lends itself primarily to the development of special HNL triggers [54,71], and underlines the importance of the accurate determination of the position of the HNL decay, as well as the HNL energy and position relative to the USER frame.

In Fig. 13, we show the distribution of the delay for HNL arrival at the MINERvA detector, compared to the arrival time of Standard Model neutrinos. Each slice of a circular plot corresponds to one bin of delay, i.e., the first bin signifies delay $\Delta t \in [0, 10]$ ns, and so on. The span of the timing bins covers the range $\Delta t \in [0, 1.6]$ μ s, which is roughly the length of one beam spill in NuMI. The left panels (a–c) are filled with the proportion of HNL events in each bin over all HNL events, whereas the right panels (d–f) are filled with the HNL events, weighted for the overall probability that each HNL would be produced, propagated, and decayed inside the detector, normalized to all HNL. Note that the radial direction is on a logarithmic scale, and that bins are only drawn if their content is greater than $10^{-3.5}$. The mass M_{N4} increases from 25 MeV/ c^2 on the top row to 250 MeV/ c^2 on the bottom row.

As expected, heavier HNL travel slower than lighter ones, and end up arriving at the detector appreciably later. Though there exist a small handful of HNL events with great delays that could end up delayed by about the length of one NuMI beam spill, they are less likely to survive long enough to reach the detector. This explains why the latest bins seem to “drop off” in the weighted right column of plots.

We also see that a significant proportion (about 10%) of HNL have a delay within the small delay bins $\sim \mathcal{O}(10)$ ns. This implies that, for detectors that support sufficient triggering sensitivity to beam-bucket timing (ns scale), it is possible in principle to obtain a trigger for delayed HNL that arrive after the Standard Model neutrinos from the beam have traversed the detector. Work utilizing such a trigger has already been done by the MicroBooNE Collaboration [53], and studied for the SBN program [71].

The simulation, much like Standard Model GENIE output, returns `EventRecords` that summarize the HNL decay. The defining features of the event record are

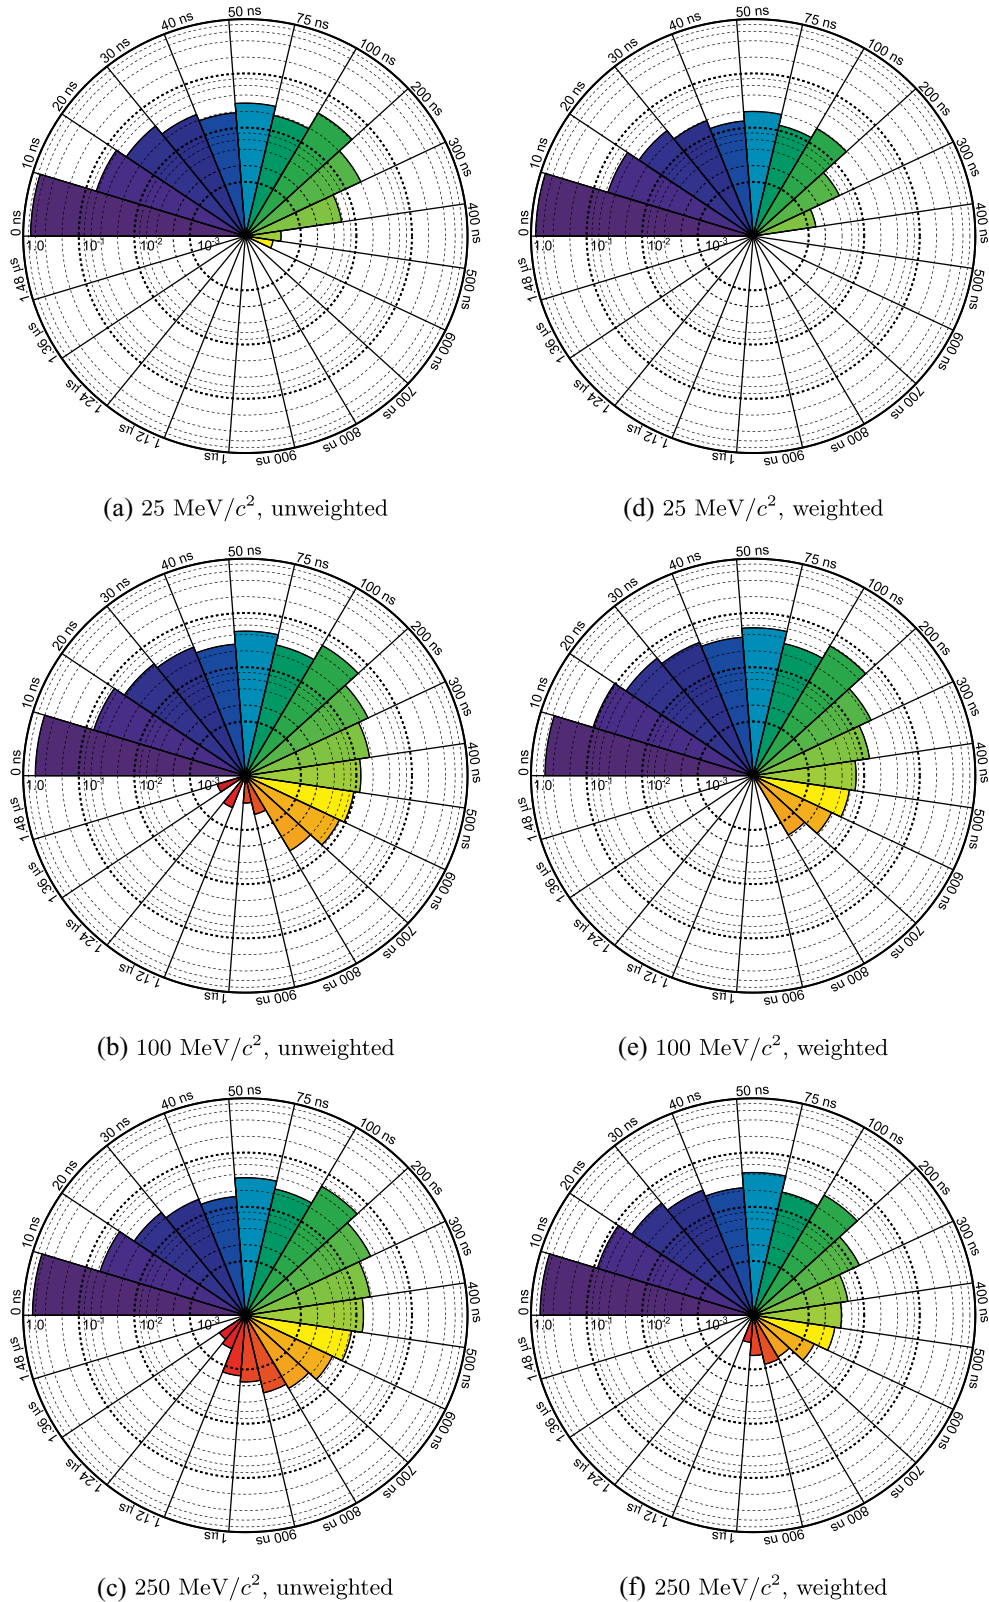


FIG. 13. Delays of HNL with respect to a SM neutrino produced by the same parent, for the MINER ν A detector. Each circular slice represents one timing bin. The final timing bin is $[1.48, 1.6] \mu\text{s}$. Panels (a–c) HNL decay events as a proportion of all simulated events. Panels (d–f) HNL decay events, weighted for the probability of production, propagation, and decay, as a proportion of the sum of weights of all events. Note the logarithmic scale on the radial axis.

- (i) The particle stack, containing each particle's PDG code and four-momentum.
- (ii) The event vertex, in USER coordinates, with the time component measuring the delay $\Delta t := t(\text{HNL reaches V}) - t(\text{SM } \nu \text{ reaches V})$.
- (iii) The weight, containing the calculated N_{POT} for this signal event.
- (iv) The "event probability," containing the HNL lifetime $\tau = 1/\Gamma_{\text{tot}}$ in units of GeV^{-1} .

We present more information about configuration and running of the module in Appendix A 1.

V. DISCUSSION

The BeamHNL module currently implements the modified neutrino kinematics including the collimation effect, polarization effects, produces realistic spatial and time distributions of HNL decay vertices, and robustly calculates production/decay rates. This corresponds to a detailed picture of the most prominent physics effects relevant for HNL below the kaon mass in all stages from parent spectra to final state distributions. Despite the wealth of physics effects already presented, however, there exist avenues for extension of the functionality of the module. Each of the following additions to the codebase would represent a successively more complete and general description of the physics of massive neutrinos, which can provide the basis for an expansion of the validity of the code from the medium beam energy $\mathcal{O}(100 \text{ MeV} - 10 \text{ GeV})$ range to the collider regime, with support for an increasing number of production channels and decay models that can be tested, such as the one described in [143], or to the atmospheric regime. For example, it would be possible in the short term to proceed in a manner similar to [101,111] to implement a spectrum of D_s mesons and τ and their decays into HNL, unlocking a higher mass range still probe-able in the medium-energy regime.

Equally, an appealing extension to the module's capability would be to generate HNL from the decays of hadrons in the atmosphere (or, generally, *not* traveling along some well-defined axis). This would open the possibility to probe HNL of extraterrestrial origin, which raises the possibility of leveraging atmospheric neutrinos to obtain higher coverage of the parameter space at low HNL masses. Work towards estimating the sensitivity of Super-Kamiokande (SK) has already been undertaken in [58,134]; as an example of a natural next step, we envisage adapting the flux description provided in that work to work with our module in order to simulate HNL decays in SK and future neutrino experiments such as DUNE, Hyper-Kamiokande or JUNO. Atmospheric experiments may also be sensitive to HNL produced through active neutrino upscattering in the Earth; IceCube is also sensitive to such a production mechanism for HNL [144].

HNL may also be produced in nuclear beta decays (for example, in the Sun through ${}^8\text{B} \rightarrow {}^8\text{Be} + e^+ + N_4$), which

raises the possibility of probing the low-mass range through direct searches. On the solar neutrino front, the Borexino Collaboration has published results of a search for the decay $N_4 \rightarrow \nu + e^+ + e^-$ [145]. This paradigm also applies to reactor HNL, where there exist limits from HNL decays in detectors (including the radiative decay $N_4 \rightarrow \nu_i + \gamma$) [146,147]. As opposed to the collider or atmospheric paradigms, though, HNL from beta decays constitute a different HNL production mechanism entirely and work to incorporate this mode into our module would necessarily be more profound. It is thus alluring to imagine the likelihood of multiple HNL production mechanisms being incorporated for a fuller description of the physics of HNL in the mass-mixing portal.

Alternatively, there has been quite some activity recently to estimate the full differential distributions of HNL decays to three bodies, including in the case for Majorana HNL [71,137,139]. The full description of higher-multiplicity polarization is an attractive goal for future high-precision experiments, assuming HNL are discovered; in principle, constructing a beam of low-energy HNL with small Lorentz boosts could allow for a measurement of the angular distributions of decay products [139], which could lend insight into the nature of the neutrino as a Dirac or a Majorana particle [137]. However, as can be seen in Fig. 11, the details of such an implementation become important, for an accelerator neutrino context, only at sub-degree detector angular resolutions.

Notwithstanding the model-independent effect of HNL polarized decay, we must point out that our current treatment of HNL decays to visible particles is dependent on the effective field theory presented in [104]. The production philosophy of parent decay \rightarrow HNL propagation \rightarrow HNL decay, on the other hand, is model independent, which motivates an effort to generalize the interface such that alternative HNL decay models can be accommodated. In our view, this should be as general as the choice of input flux simulation or geometry description, to enable versatile usage of this module for different theories that differ only as to the decay widths they predict. Given that the most stringent limits on the mixing elements $\{|U_{\alpha 4}|^2\}$ are placed by dedicated experimental searches that look for model-dependent kinematic signatures [148], it is important to enable the implementation of models that all share the same conception of HNL production and decay, in order to maximize the utility of experimental data.

The HNL parameter space under the flavor mixing assumptions $|U_{e4}|^2 : |U_{\mu 4}|^2 : |U_{\tau 4}|^2 = 1 : 0 : 0, 0 : 1 : 0$ is constrained by particle physics experiments (see, for example, [67,78] and references therein). Complementary bounds are obtained by measuring the primordial elemental abundances that result from Big Bang nucleosynthesis (BBN); if the lifetime τ of HNL is long enough, they decay to mesons in the primordial plasma, which in turn forces proton and neutron number densities towards equilibrium and alters




FIG. 14. Parameter space limits for HNL searches below the kaon mass, with limits used from [57,62,63,67,149,150] and BBN contours shown at $\tau = 0.023, 0.1$ (filled), 0.5 s [152].

the primordial elemental abundances. In Fig. 14, we have shown in panels (a) and (b) the currently excluded parameter space, with the colored region corresponding to experimental limits [57,62,63,67,149,150], and the shaded gray area corresponding to a BBN constraint $\tau \leq 0.1$ s [151]. Also shown are the constraints for $\tau \leq 0.023$ s [152], ≤ 0.5 s as solid gray lines.

As can be seen, for these mixing assumptions the conjunction of experimental limits with BBN almost completely covers the available parameter space for a single HNL below the kaon mass. Note, however, that the constraint $\tau \lesssim \tau_0$, for τ_0 some suitable upper bound, places constraints on the *total* effective mixing $\sum_{\alpha} |U_{\alpha 4}|^2$, which makes the curves in Fig. 14 dependent on the assumed flavor structure. In panel (c), we show how assuming a large enough mixing with τ modifies the BBN bounds; assuming $|U_{e4}|^2 = |U_{\mu 4}|^2 = |U_{\alpha}|^2$ and $|U_{\tau 4}|^2 = 10^{-7}$, the parameter space above $(M_{N4}, |U_{\alpha 4}|^2) \sim (200 \text{ MeV}/c^2, 10^{-10})$ is completely free of BBN constraints for the same HNL lifetimes. We show no experimental bounds in panel (c), as these are extracted under a particular flavor mixing assumption; to the best of our knowledge, no search has explored this particular mixing scenario.

A final consideration is the implication of the realistic picture where multiple HNL can oscillate into each other; it was shown in [153] that this can significantly modify the decay rates into visible channels. There has been some discussion in [153–156] of what the implications of multiple HNL would be on the current most strict bounds on sterile-neutrino mixing into light neutrinos. Specifically, if there exist multiple HNL ($n > 1$, which is required by realistic extensions to the Standard Model that explain neutrino masses), then the current most stringent limits on the HNL parameter space may be significantly relaxed. This occurs because, if there exists at least one other HNL N_5 such that $\Delta M_{45} \ll M_{N_{4,5}}$, then the decay rates to visible final states can change drastically, weakening the current exclusion bounds that are made under the assumption of one heavy neutrino state.

It is, however, possible that the experimental bounds could be relaxed in the multiple-HNL picture, once again opening the allowed window for HNL at $M_{N4} < m_K$ (for instance, [154] showed that ATLAS bounds on HNL are relaxed by up to 2 orders of magnitude, depending on the prevalence of lepton-number violating processes and the precise ratios of the mixing elements). A numerical study showing how the parameter space with two HNL is still compatible with both experimental and BBN constraints was carried out in [155]. Future accelerator neutrino experiments will be capable of exploring this parameter space region, which emphasizes the crucial importance of consistent, well-founded modeling for HNL searches, as well as the desirability of a multiple-HNL implementation to support these efforts.

VI. SUMMARY

Heavy neutral leptons are among the most natural and minimal beyond the Standard Model theories, which are motivated by the measurement of nonzero neutrino masses, the first evidence that the Standard Model is incomplete. They have very important implications for both particle physics and cosmology, which makes them prime targets for searches in accelerator neutrino experiments and beyond. Already at masses $\mathcal{O}(100 \text{ MeV}/c^2)$, there is a rich phenomenology derived from the massive neutrino kinematics that is qualitatively different from massless Standard Model neutrinos. Any simulation treating HNL needs to carefully implement this phenomenology. Indeed, a definite detection of HNL would have ground-breaking implications; on top of being an incontrovertible piece of evidence of beyond the Standard Model physics and a probe for the details of neutrino mass generation, knowledge of how to produce HNL would open up the landscape for precision searches of HNL decays and, through this, the possibility for determining the neutrino’s nature.

Our goal is to provide a single, sophisticated simulation that incorporates this new physics in a self-consistent manner, for use with experiments past, present, and future.

This paper has reviewed the theory of HNL production from heavy meson decays and presented the new physics effects associated with neutrino mass, comparing to the Standard Model case of massless neutrinos. We also commented on extant similar simulation efforts, showing how this work is complementary to the alternative production model of HNL by upscattering and describing the benefits associated with accepting a generalized hadron beamline simulation and detector description as inputs. Importantly, we have decoupled the description of the primary beam and the decay volume from the intrinsic physics of the massive neutrino in a general manner, which we presented here. We then discussed interesting phenomenology that is crucial to properly simulate HNL from particle decay. We described the philosophy and crucial ingredients of the full modeling chain, from production and propagation to decay in a detector.

With our model, we have constructed a fully factorizable interface to beamline simulation and detector geometry that simulates the detailed phenomenology of HNL in a beamline. The module can automatically incorporate future developments in beamline modeling, as these will generally be contained in the beamline simulation provided as input. Likewise, as future experiments make design choices about what detector setups to use, this module allows them to study in detail the prospects for HNL detection for all the potential detectors being considered.

Though this implementation leverages the GENIE framework, the philosophy behind it is general; likewise, though the first use case could be the FNAL experiments, other accelerator neutrino contexts could also well use this module. Furthermore, we sketch a possible development strategy for extending the usability of this already robust simulation, in preparation for exciting future experiments.

ACKNOWLEDGMENTS

We would like to thank C. Andreopoulos, R. Hatcher, and M. Roda for valuable input and guidance during the GENIE development process. We further thank L. Fields for helpful discussions on simulating the neutrino flux. We also thank R. Shrock for valuable comments on the manuscript. X.L. is supported by the STFC (UK) Grants No. ST/S003533/1 and No. ST/S003533/2.

APPENDIX: CODE AND ALGORITHM DETAILS

1. Quick guide to running BeamHNL in GENIE

A fork of GENIE with the BeamHNL module implementation may be found on GitHub [157]. The configuration file has been designed with ease of use in batch jobs, such as might be submitted to a computing grid, in mind. For this reason, there exists a single file `config/CommonHNL.xml` that houses all the options that the module's components need to know about in order to simulate the HNL production, propagation, and decay. A

sample file may be found in `src/contrib/beamhnl`, and explanations for the various options are available in `config/NHLPrimaryVtxGenerator.xml`. In this section, we will briefly outline the main components.

- (i) `ParameterSpace` defines M_{N_4} in GeV/c^2 , $\{|U_{\alpha 4}|^2\}$, and Dirac vs Majorana nature of the HNL.
- (ii) `InterestingChannels` enumerates the 10 decay channels (the 7 summarized in Table II and $N_4 \rightarrow \pi^\pm \pi^0 \ell^\mp, \pi^0 \pi^0 \nu$) kinematically accessible to an HNL below the kaon mass. Entries set to `false` will be inhibited from entering the event record, whereas `true` entries are treated as valid truth signal channels.
- (iii) `CoordinateXForm` defines the unique translation and rotation vectors \mathbf{T}, \mathbf{R} from NEAR to BEAM and from NEAR to USER coordinates. \mathbf{T} is given in metres, with respect to the NEAR system, and \mathbf{R} is a vector of 3 Euler angles, following the “extrinsic $X - Z - X$ ” convention; that is, the rotation matrix R is given as

$$R(\alpha, \beta, \gamma) = R_X(\alpha)R_Z(\beta)R_X(\gamma), \quad (\text{A1})$$

where X, Y, Z are the fixed NEAR axes.

- (iv) `FluxCalc` provides switches for the user to enable/disable certain features; namely, the module's polarization accept/reject weight, and whether the simulation should evaluate Eq. (A15) assuming the parent to be perfectly focused (setting ζ_+ to one-half the detector's angular opening, and $\zeta_- = 0$)

For the purposes of running this module, an input `dk2nu`-like flux ROOT flat-tree and a ROOT geometry file describing the detector are required. Sample inputs that are module compliant are supplied along with the source code in `src/contrib/beamhnl` for the user to be able to run the module immediately.

A folder `flatDk2nus` contains two flat ROOT trees, corresponding to one `dk2nu` flux file from NuMI in neutrino mode, and one in antineutrino mode [132]. It also contains scripts to produce ROOT flat trees from `dk2nu` flux files. These flat trees can then be used as input of the BeamHNL module. Detailed instructions on how to generate these are written in the README file located inside the folder.

Three ROOT macros `makeBox.C`, `makeCylinder.C`, and `makeHexagon.C` along with three respective outputs are also inside the `contrib/beamhnl` directory, which allow the user to make three different ROOT geometries of arbitrary dimensions and rotation with respect to the USER coordinate system.

To enable the BeamHNL module, the user must first configure GENIE appropriately by adding the following line to the `configure` script (an example can be seen at the GENIE website [158]):




FIG. 15. Kinematic scaling factor and $\mathfrak{B}_{\text{channel}}$ for the HNL production channels resulting from π^\pm decay. See text for details.

-- enable -- heavy -- neutral -- lepton \

After running make, the BeamHNL module is ready. The archetypal run command is

```
gevgen_hnl -n <nSignalEvents> -f <path/to/flux/dir> -g <path/to/geom/file.root>
```

Detailed instructions can be referred to by passing the `-h` flag.

2. HNL production and decay rates

The available production channels for an HNL with mass $M_{N_4} \lesssim m_K$ are summarized in Table I.³ For each channel, the threshold, kinematic scaling \mathcal{K} and SM branching ratio are listed. The kinematic scaling \mathcal{K} is defined through the decay width as in Eq. (5),

$$\Gamma(P \rightarrow N_4 + \ell_\alpha + \dots) = |U_{\alpha 4}|^2 \mathcal{K} \cdot \Gamma(P \rightarrow \nu + \ell_\alpha + \dots). \quad (\text{A2})$$

The decay width is then, for a given channel, written as

$$\Gamma(M_{N_4}, |U_{\alpha 4}|^2) = \mathcal{K}(M_{N_4}) \cdot |U_{\alpha 4}|^2, \quad (\text{A3})$$

³Throughout this section, we will engage in a mild abuse of notation and write down such expressions as $K^\pm \rightarrow N_4 + \mu^\pm$. The reader will notice that N_4 has been defined in Eq. (4) as a *mass eigenstate*, rather than a *flavor eigenstate*. In order to shorten notation, we will sacrifice formal correctness (i.e., $K^\pm \rightarrow \nu_\mu + \mu^\pm, \nu_\mu = \sum_i U_{\mu i} \nu_i + U_{\mu 4} N_4$) and directly concentrate on the mixed-in element $U_{\alpha 4} N_4$, hoping that the implied association does not cause confusion.

For each parent species ($\pi^\pm, K^\pm, K_L^0, \mu^\pm$) a series of scores $\{s_i\}, s_{i+1} = s_i + \Delta s$ is constructed, which is then used to determine the production channel for the HNL. The scores are calculated as

$$\Delta s = \frac{\mathfrak{B}_{\text{channel}}}{\mathfrak{B}_{\text{tot}}}, \quad (\text{A4})$$

where $\mathfrak{B}_{\text{tot}}$ is the sum of $\mathfrak{B}_{\text{channel}}$ over all kinematically accessible HNL production channels (see Table II).

We illustrate these steps for HNL production by pion decay, with the help of Fig. 15. First, we calculate the kinematic scaling factor $\mathcal{K}_{\text{channel}}$ according to Table I (shown in the left panel). Afterwards, we multiply each by the SM branching ratio ($\mathfrak{B}_{\pi \rightarrow e}^{(\text{SM})} = 1.23 \times 10^{-4}$, $\mathfrak{B}_{\pi \rightarrow \mu}^{(\text{SM})} = 0.999877$) to obtain the $\mathfrak{B}_{\text{channel}}$ (shown in the right panel).

The following definitions are used for the case of 2-body HNL production [110]:

$$\delta(m, M) \equiv \delta_M^m := \frac{m^2}{M^2}, \quad (\text{A5a})$$

$$\lambda(x, y, z) := x^2 + y^2 + z^2 - 2(xy + yz + zx), \quad (\text{A5b})$$

$$f_m(x, y) := x + y - (x - y)^2, \quad (\text{A5c})$$

$$\rho(x, y) := f_m(x, y) \cdot \lambda^{1/2}(x, y, 1), \quad (\text{A5d})$$

$$\mathcal{P}_\ell(x, y) := \frac{\rho(x, y)}{x(1-x)^2}. \quad (\text{A5e})$$

Calculations of the three-body HNL scaling factor for production have been done in the literature. We have used the helicity-summed scaling factors reported in [111] to construct through interpolation a scaling function $\mathcal{S}_{P\ell}$, with $P = K^\pm, K_L^0$, and $\ell = e, \mu$. For the special case of muon decay to HNL, we start from the known decay $\mu^\pm \rightarrow \nu_\mu \nu_e e^\pm$ and then “promote” either the ν_μ or ν_e to HNL, depending on the mixings $|U_{e4}|^2, |U_{\mu 4}|^2$.

We adopt the HNL decay widths in the context of an effective field theory describing interactions of HNL with mesons, as detailed in [104]. For the double-pion channels $N_4 \rightarrow \pi^\pm \pi^0 \ell^\mp, \pi^0 \pi^0 \nu$, whose thresholds for emission lie below the kaon mass, the decay widths in the literature have been noted to be dominated by the chain involving the emission of an on-shell ρ . Because $m_\rho > m_K$, this argument cannot be applied in the case $M_{N_4} < m_K$, and it is necessary to perform an explicit calculation to estimate the decay width resulting from double-pion channels. To that effect, we have used the Lagrangian for Dirac particles, made public by the authors of [104] in the FeynRules model database [118], and extracted the double-differential decay rate over final-state particle energies using FeynRules [112], FeynArts [113], and FeynCalc [114–116]. The resulting expression for the two-dimensional $d^2\Gamma/dE_\pi dE_\ell$ depends on the energy of the π^0 due to energy conservation in the HNL rest frame. Therefore, the integrated decay rate is obtained in runtime by numerically integrating the differential decay rate.

As can be seen from the relevant tree-level Feynman diagrams in Fig. 16, the decay $N_4 \rightarrow \pi^\pm \pi^0 \ell^\mp$ is mediated by both light and heavy neutrinos. Compared




FIG. 16. Feynman diagrams for the decay $N_4 \rightarrow \pi^\pm + \pi^0 + \ell^\mp$.

to the left-hand diagrams, the right-hand diagram where N_4 enters both pion vertices is suppressed by a factor $|U_{\alpha 4}|^2/|U_{\alpha i}|^2 \ll 1$, and is safely ignored.

For the $\pi^0 \pi^0 \nu$ diagrams, the intermediate propagator is always a N_4 , which renders the entire channel’s decay width subleading to $\pi^\pm \pi^0 \ell^\mp$.

Comparing $N_4 \rightarrow \pi^\pm \pi^0 \ell_\alpha^\pm$, the $\alpha = e$ case has a larger decay width due to the larger phase space available to the final state. The result is

$$\frac{\Gamma(N_4 \rightarrow \pi^+ \pi^0 e^-)}{\sum_{\text{all lighter channels}} \Gamma} \lesssim 10^{-11}, \quad (\text{A6})$$

confirming the argument that the essential process for double-pion production is decays of HNL into on-shell ρ . The calculation proceeds similarly for Majorana HNL.

To conserve computing time, we only simulate those channels that the user has explicitly defined in the configuration file as being “interesting” as a signal. For each run of the module, a C++ `std::map<HNLDecayMode_t, double>` is constructed that contains the accessible and interesting channels and their decay widths. These are then used to construct scores similarly to Eq. (A4), after which a standard Monte Carlo transformation method [159] maps a uniform random number to a decay channel and fills the appropriate decay product list for use with GENIE’s phase-space generator. We note that the HNL lifetime is computed from the full list of all kinematically available channels, regardless of whether these channels are interesting or not.

To calculate the decay rates, we make use of the following definitions found in [104]:

$$B_1 := \frac{1}{4}(1 - 4 \sin^2 \theta_W + 8 \sin^4 \theta_W), \quad (\text{A7a})$$

$$B_2 := \frac{1}{2} \sin^2 \theta_W (2 \sin^2 \theta_W - 1), \quad (\text{A7b})$$

$$\begin{aligned} C_\alpha(M_{N_4}) := & \sum_\alpha |U_{\alpha 4}|^2 \cdot [F_1(m_\alpha/M_{N_4})B_1 \\ & + F_2(m_\alpha/M_{N_4})B_2], \end{aligned} \quad (\text{A7c})$$

$$\begin{aligned} \mathcal{D}_\alpha(M_{N_4}) := & |U_{\alpha 4}|^2 \sin^2 \theta_W \\ & \times [2F_1(M_{N_4}) + F_2(M_{N_4})], \end{aligned} \quad (\text{A7d})$$

$$\begin{aligned} \mathcal{G}(x, y) := & \lambda^{1/2}(x, y, 1) \\ & \times [1 - y^2 - x^2(2 - x^2 + y^2)], \end{aligned} \quad (\text{A7e})$$

$$\begin{aligned} F_1(x) = & (1 - 14x^2 - 2x^4 - 12x^6) \sqrt{1 - 4x^2} \\ & + 12x^4(x^4 - 1)L(x), \end{aligned} \quad (\text{A7f})$$




FIG. 17. Calculation of deviation angles ζ_{\mp} . The parent's momentum \mathbf{p}_P is projected to the point V_0 such that $z_{V_0} = z_C$, with C the center of the detector. The entry and exit points V_{\mp} lie on the line $\epsilon: \mathbf{r}(u) = \mathbf{r}_{V_0} + u \cdot \boldsymbol{\delta}$, where $\boldsymbol{\delta}$ is a sweep direction: $\boldsymbol{\delta} := \mathbf{r}_C - \mathbf{r}_{V_0}$. The angles ζ_{\mp} are $\langle \mathbf{r}_{\mp}, \mathbf{p}_P \rangle$.




FIG. 18. Extrinsic Euler rotations. The fixed system (X, Y, Z) is rotated to the new system (X_U, Y_U, Z_U) by applying successive rotations first about X (green), then about Z (blue), and finally about X again (red).

$$F_2(x) = 4x^2(2 + 10x^2 - 12x^4)\sqrt{1 - 4x^2} + 24x^4(1 - 2x^2 + 2x^4)L(x), \quad (\text{A7g})$$

$$L(x) = \ln \left[\frac{1 - 3x^2 - (1 - x^2)\sqrt{1 - 4x^2}}{x^2(1 + \sqrt{1 - 4x^2})} \right], \quad (\text{A7h})$$

where θ_W is the Weinberg mixing angle, and $\lambda(x, y, z)$ is the Källén function defined in Eq. (A5b).

3. Coordinate systems

A general detector can be both displaced and rotated arbitrarily with respect to the NEAR frame. One can parametrize any such configuration by two vectors: one translation and one rotation. We have made the choice to use extrinsic Euler angles to write rotations (see Fig. 18): that is, the rotation matrix describing the transition between two coordinate systems (X, Y, Z) , (X_U, Y_U, Z_U) centered around the same point is written as

$$\begin{aligned} R(\alpha, \beta, \gamma) &= R_X(\gamma)R_Z(\beta)R_X(\alpha) \\ &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_\gamma & -s_\gamma \\ 0 & s_\gamma & c_\gamma \end{pmatrix} \begin{pmatrix} c_\beta & -s_\beta & 0 \\ s_\beta & c_\beta & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_\alpha & -s_\alpha \\ 0 & s_\alpha & c_\alpha \end{pmatrix} \\ &= \begin{pmatrix} c_\beta & -c_\alpha s_\beta & s_\alpha s_\beta \\ s_\beta c_\gamma & c_\alpha c_\beta c_\gamma - s_\alpha s_\gamma & -s_\alpha c_\beta c_\gamma - c_\alpha s_\gamma \\ s_\beta s_\gamma & s_\alpha c_\gamma + c_\alpha c_\beta s_\gamma & c_\alpha c_\gamma - s_\alpha c_\beta s_\gamma \end{pmatrix}, \quad (\text{A8}) \end{aligned}$$

using $c_\theta, s_\theta \equiv \cos \theta, \sin \theta$.

4. Bookkeeping

There are a certain number of quantities that must be kept track of during the simulation in order to correctly estimate the number of signal events, given a detector volume and number of protons on target.⁴ These can be summarized in Fig. 19; we shall explain the steps taken forthwith.

Suppose that the user has selected a channel C as the channel of interest for a particular detector; for example, $N_4 \rightarrow \mu^- \pi^+$. We work our way backwards in order to obtain, *for each signal event*, an estimate of the number of POT N_{POT} that would result in one signal event occurring in the detector, and return this estimate as a weight attached to the `EventRecord` describing the signal event, which makes POT counting a straightforward loop on the analysis side.

First to be obtained is the expected number of *total* HNL decays occurring in the detector, signal or not (for example, invisible decays $N_4 \rightarrow \nu \nu \nu$ are almost never going to be considered signal events due to the inability to detect neutrinos in the final state directly), as




FIG. 19. Sequence for estimating number of POT for each event, working backwards. See text for definitions.

$$N_H = \frac{\sum_{i \in \text{all channels}} \Gamma_i}{\Gamma_C} \equiv \frac{\Gamma_{\text{tot}}}{\Gamma_C}. \quad (\text{A9})$$

One then takes the detector geometry into account, by considering the HNL's lifetime $\tau = \hbar/\Gamma_{\text{tot}}$ and requiring that the HNL decays inside the detector volume. Suppose a beam of HNL of rest-frame lifetime τ and velocity βc propagates along the z axis, and the detector volume is in between the planes $z = z_1$ and $z = z_2$. The probability distribution for the HNL decay location is

$$p(z) = \frac{1}{\tilde{N}} \exp\left(-\frac{z}{\beta c \gamma \tau}\right), \quad (\text{A10})$$

where \tilde{N} is some dimensionful normalization constant. The probability of decay inside the detector is then

$$P(z_1 \leq z_{\text{decay}} \leq z_2) = \int_{z_1}^{z_2} du p(u), \quad (\text{A11})$$

which yields

$$\begin{aligned} P &= \frac{\beta \gamma c \tau}{\tilde{N}} \exp\left(-\frac{z_1}{\beta \gamma c \tau}\right) \left[1 - \exp\left(-\frac{z_2 - z_1}{\beta \gamma c \tau}\right)\right] \\ &= P(\text{arrival}) \cdot P(\text{decay}|\text{arrival}). \end{aligned} \quad (\text{A12})$$

Equation (A12) states that the HNL must first survive long enough to reach the detector, and then decay promptly while inside it.

We can generalize this to the full 3D picture. Let D , E , and X be the HNL production point, and the intersections of its trajectory with the detector at entry and exit, respectively. Then for every N_P HNL emitted that could be accepted, N_A will survive until the detector, and N_H will decay without exiting the detector. These quantities are straightforwardly obtained:

$$N_A = N_H \cdot \exp\left(\frac{\ell}{\beta c \gamma \tau}\right), \quad (\text{A13})$$

$$N_P = N_A \cdot \exp\left(\frac{L}{\beta c \gamma \tau}\right), \quad (\text{A14})$$

where ℓ, L are the distances $|\mathbf{r}_X - \mathbf{r}_E|, |\mathbf{r}_D - \mathbf{r}_E|$.

Next, we factor in those HNL that decayed before reaching the detector volume to get the total number N_P

⁴By ‘‘protons on target’’ we will account for the mean number of POT for a single HNL decay to *signal* event. This information is calculated taking the input detector volume and beamline simulation; different inputs will yield different outputs.

of HNL produced in the accepted region, using once again Eq. (A13) but substituting $\ell = |\mathbf{r}_X - \mathbf{r}_E|$ for $L = |\mathbf{r}_E - \mathbf{r}_D|$.

Further backward, we have the estimate of acceptance; only those HNL emitted in the correct angular region would intersect the detector. It is assumed that the decays of parents are isotropic in the parents' rest frame, which is correct for pseudoscalar mesons (π^\pm, K^\pm, K^0). Furthermore, the detector is assumed to be sufficiently far away that the small angle approximation $\sin \theta \simeq \theta$ is valid. The size and position of the detector defines a window in the observer's frame, which is then transformed into a rest-frame window using the collimation-effect function $f: [0, \pi] \rightarrow [0, \pi]$ described in Fig. 4. The lab-frame emission angle can be interpreted as an angular deviation of the HNL's trajectory from the parent's momentum. The angular size of the detector is given by $\zeta_+ - \zeta_-$, where the angles ζ_\mp are the minimum and maximum deviation angles for which the HNL's trajectory can intersect the detector. Figure 17 shows how these deviation angles are calculated. A ‘‘sweep direction’’ δ is constructed from the parent's momentum \mathbf{p}_P and the detector center C , and the points of entry V_- and exit V_+ along this sweep are obtained, giving the deviation angles as the angles between the parent momentum and the vectors $\mathbf{r}_-, \mathbf{r}_+$.

We thus estimate the acceptance correction \mathcal{A} induced by the collimation effect of HNL becoming dominated by their parent's Lorentz boost, as

$$\mathcal{A} = \frac{|\mathcal{I}_{N4}|}{|\mathcal{I}_\nu|}, \quad (\text{A15})$$

where \mathcal{I} is the pre-image under f of the angular opening $[\zeta_-, \zeta_+]$. For large HNL masses, it is $\mathcal{I}_{N4} = \mathcal{I}_F \sqcup \mathcal{I}_B$ where $\mathcal{I}_{F,B}$ are the forward and backward rest-frame angular regions where HNL can be accepted. One can contrast this with the case of light neutrinos, where f increases monotonously and only forward emitted neutrinos can reach the detector. In the end, the number of total HNL emitted (‘‘ancestry events’’) is given by

$$N_C = \frac{N_P}{\sum_{\alpha=e,\mu,\tau} |U_{\alpha 4}|^2} \cdot \frac{1}{\omega_{\text{det}} \mathcal{B}^2 \mathcal{A}}, \quad (\text{A16})$$

where \mathcal{B}, \mathcal{A} were defined in Eqs. (11) and (A15), and ω_{det} is the angular size of the detector in the lab frame, with D at the origin. The angular size of the detector in the parent's rest frame is then $\omega_{\text{det}} \mathcal{B}^2$, assuming its face is roughly perpendicular to the parent momentum; the acceptance correction \mathcal{A} parametrizes the intrinsic increase of the probability that a randomly chosen direction for HNL emission will be boosted such that the HNL gets accepted. The prefactor $(\sum_\alpha |U_{\alpha 4}|^2)^{-1}$ is inserted to correct for the fact that, to conserve computing resources, we assume all parent decays result in an HNL. Finally, we estimate

$N_{\text{POT}} = N_C \cdot \tilde{N}(P)$, $\tilde{N}(P) = n(\text{all particles})/n(P)$, where n is the number of particles produced in a $p + \text{target}$ interaction. For example, if in $n(\pi^+) = 0.4$, $n(K^+) = 0.1$, $n(K^0) \simeq n(K_L^0) = 0.05$, $n(p) = 0.15$, $n(\mu^+ + \text{other}) \simeq 0.05$, then the relevant factors are

$$\tilde{N}\{\pi^+, K^+, K_L^0, \mu^+, p\} = \{1.875, 7.5, 15, 5, 15\}. \quad (\text{A17})$$

For practical purposes, the code expects the values $\tilde{N}(P)$ as inputs from the user in the configuration file.

-
- [1] B. Pontecorvo, *Sov. Phys. JETP* **33**, 429 (1958).
 [2] B. Pontecorvo, *Sov. Phys. JETP* **26**, 984 (1968).
 [3] S. Bilenky, *Nucl. Phys.* **B908**, 2 (2016).
 [4] R. Davis, D. S. Harmer, and K. C. Hoffman, *Phys. Rev. Lett.* **20**, 1205 (1968).
 [5] Q. R. Ahmad *et al.* (SNO Collaboration), *Phys. Rev. Lett.* **89**, 011301 (2002).
 [6] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. Lett.* **81**, 1562 (1998).
 [7] K. Eguchi *et al.* (KamLAND Collaboration), *Phys. Rev. Lett.* **90**, 021802 (2003).
 [8] P. F. de Salas, D. V. Forero, S. Gariazzo, P. Martínez-Miravé, O. Mena, C. A. Ternes, M. Tórtola, and J. W. F. Valle, *J. High Energy Phys.* **02** (2021) 71.
 [9] W. Hampel *et al.* (GALLEX Collaboration), *Phys. Lett. B* **447**, 127 (1999).
 [10] M. Altmann *et al.* (GNO Collaboration), *Phys. Lett. B* **616**, 174 (2005).
 [11] Y. Fukuda *et al.*, *Phys. Rev. Lett.* **77**, 1683 (1996).
 [12] M. Agostini *et al.* (Borexino Collaboration), *Nature (London)* **587**, 577 (2020).
 [13] M. H. Ahn *et al.* (K2K Collaboration), *Phys. Rev. D* **74**, 072003 (2006).
 [14] P. Adamson *et al.*, *Phys. Rev. Lett.* **112**, 191801 (2014).
 [15] N. Agafonova *et al.* (OPERA Collaboration), *Phys. Rev. Lett.* **120**, 211801 (2018).
 [16] A. Aguilar *et al.* (LSND Collaboration), *Phys. Rev. D* **64**, 112007 (2001).
 [17] B. Armbruster *et al.* (KARMEN Collaboration), *Phys. Rev. D* **65**, 112001 (2002).
 [18] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Phys. Rev. D* **103**, 052002 (2021).
 [19] M. Apollonio *et al.*, *Eur. Phys. J. C* **27**, 331 (2003).
 [20] F. Boehm *et al.*, *Phys. Rev. D* **64**, 112001 (2001).
 [21] P. Adamson *et al.* (MINOS+ Collaboration), *Phys. Rev. Lett.* **125**, 131802 (2020).
 [22] K. Abe *et al.* (T2K Collaboration), *Phys. Rev. D* **103**, 112008 (2021).
 [23] M. A. Acero *et al.* (NOvA Collaboration), *Phys. Rev. Lett.* **123**, 151803 (2019).
 [24] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. Lett.* **128**, 241801 (2022).
 [25] D. Adey *et al.* (Daya Bay Collaboration), *Phys. Rev. D* **100**, 052004 (2019).
 [26] G. Bak *et al.* (RENO Collaboration), *Phys. Rev. Lett.* **121**, 201801 (2018).
 [27] H. de Kerret *et al.* (Double Chooz Collaboration), *Nat. Phys.* **16**, 558 (2020).
 [28] P. A. N. Machado, O. Palamara, and D. W. Schmitz, *Annu. Rev. Nucl. Part. Sci.* **69**, 363 (2019).
 [29] B. Abi *et al.*, arXiv:2002.03005.
 [30] F. An *et al.*, *J. Phys. G* **43**, 030401 (2016).
 [31] K. Abe *et al.* (Hyper-Kamiokande Working Group), arXiv:1412.4673.
 [32] M. Aker *et al.* (KATRIN Collaboration), *Nat. Phys.* **18**, 160 (2022).
 [33] R. L. Workman *et al.* (Particle Data Group), *Prog. Theor. Exp. Phys.* **2022**, 083C01 (2022).
 [34] J. A. Formaggio, A. L. C. de Gouvêa, and R. G. H. Robertson, *Phys. Rep.* **914**, 1 (2021).
 [35] O. Miranda and J. Valle, *Nucl. Phys.* **B908**, 436 (2016).
 [36] A. Abada, C. Biggio, F. Bonnet, M. B. Gavela, and T. Hambye, *J. High Energy Phys.* **12** (2007) 061.
 [37] Y. Cai, T. Han, T. Li, and R. Ruiz, *Front. Phys.* **6**, 40 (2018).
 [38] T. Asaka and M. Shaposhnikov, *Phys. Lett. B* **620**, 17 (2005).
 [39] I. Brivio, K. Moffat, S. Pascoli, S. T. Petcov, and J. Turner, *J. High Energy Phys.* **10** (2019) 059.
 [40] V. Brdar, A. J. Helmboldt, S. Iwamoto, and K. Schmitz, *Phys. Rev. D* **100**, 075029 (2019).
 [41] A. Boyarsky, M. Drewes, T. Lasserre, S. Mertens, and O. Ruchayskiy, *Prog. Part. Nucl. Phys.* **104**, 1 (2019).
 [42] M. Shaposhnikov, *Nucl. Phys.* **B763**, 49 (2007).
 [43] J. Bernabéu, A. Santamaria, J. Vidal, A. Mendez, and J. W. F. Valle, *Phys. Lett. B* **187**, 303 (1987).
 [44] F. F. Deppisch, P. S. Bhupal Dev, and A. Pilaftsis, *New J. Phys.* **17**, 075019 (2015).
 [45] K. N. Abazajian *et al.*, arXiv:1204.5379.
 [46] A. Ibarra, E. Molinaro, and S. T. Petcov, *Phys. Rev. D* **84**, 013005 (2011).
 [47] M. Fukugita and T. Yanagida, *Phys. Lett.* **174B**, 45 (1986).
 [48] E. K. Akhmedov, V. A. Rubakov, and A. Y. Smirnov, *Phys. Rev. Lett.* **81**, 1359 (1998).
 [49] D. Gorbunov and M. Shaposhnikov, *J. High Energy Phys.* **10** (2007) 015.
 [50] E. Bulbul, M. Markevitch, A. Foster, R. K. Smith, M. Loewenstein, and S. W. Randall, *Astrophys. J.* **789**, 13 (2014).
 [51] S. Bhargava *et al.*, *Mon. Not. R. Astron. Soc.* **497**, 656 (2020).
 [52] E. M. Silich, K. Jahoda, L. Angelini, P. Kaaret, A. Zajczyk, D. M. LaRocca, R. Ringuette, and J. Richardson, *Astrophys. J.* **916**, 2 (2021).
 [53] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **101**, 38 (2020).

- [54] S. D. Porzio, Searches for heavy neutral lepton decays in the MicroBooNE detector, Ph.D. thesis, Manchester University, 2019.
- [55] K. J. Kelly and P. A. N. Machado, *Phys. Rev. D* **104**, 055015 (2021).
- [56] O. R. Y. Goodwin, Search for Higgs portal scalars and heavy neutral leptons decaying in the MicroBooNE detector, Ph.D. thesis, Manchester University, 2022.
- [57] K. Abe *et al.* (T2K Collaboration), *Phys. Rev. D* **100**, 052006 (2019).
- [58] P. Coloma, P. Hernández, V. Muñoz, and I. M. Shoemaker, *Eur. Phys. J. C* **80**, 235 (2020).
- [59] R. Acciarri *et al.* (ArgoNeuT Collaboration), *Phys. Rev. Lett.* **127**, 121801 (2021).
- [60] R. Aaij *et al.* (LHCb Collaboration), *Eur. Phys. J. C* **81**, 248 (2021).
- [61] A. V. Artamonov *et al.* (E949 Collaboration), *Phys. Rev. D* **91**, 052001 (2015).
- [62] A. Aguilar-Arevalo *et al.* (PIENU Collaboration), *Phys. Rev. D* **97**, 072012 (2018).
- [63] A. Aguilar-Arevalo *et al.*, *Phys. Lett. B* **798**, 134980 (2019).
- [64] G. Bernardi *et al.*, *Phys. Lett.* **166B**, 479 (1986).
- [65] G. Bernardi *et al.*, *Phys. Lett. B* **203**, 332 (1988).
- [66] E. Cortina Gil *et al.* (NA62 Collaboration), *Phys. Lett. B* **778**, 137 (2018).
- [67] A. M. Abdullahi *et al.*, *J. Phys. G* **50**, 020501 (2023).
- [68] K. Abe *et al.* (Hyper-Kamiokande Proto-Collaboration), [arXiv:1805.04163](https://arxiv.org/abs/1805.04163).
- [69] B. Abi *et al.* (DUNE Collaboration), *Eur. Phys. J. C* **81**, 322 (2021).
- [70] M. Breitbach, L. Buonocore, C. Frugiuele, J. Kopp, and L. Mitnacht, *J. High Energy Phys.* 01 (2022) 048.
- [71] P. Ballett, S. Pascoli, and M. Ross-Lonergan, *J. High Energy Phys.* 04 (2017) 102.
- [72] C. Ahdida *et al.* (SHiP Collaboration), *J. High Energy Phys.* 04 (2019) 077.
- [73] D. Gorbunov, I. Krasnov, Y. Kudenko, and S. Suvorov, *Phys. Lett. B* **810**, 135817 (2020).
- [74] Z. S. Wang and K. Wang, *Phys. Rev. D* **101**, 075046 (2020).
- [75] J. L. Feng *et al.*, *J. Phys. G* **50**, 030501 (2023).
- [76] S. Cerci *et al.*, *J. High Energy Phys.* 06 (2022) 110.
- [77] F. Kling and S. Trojanowski, *Phys. Rev. D* **97**, 095016 (2018).
- [78] D. Bryman and R. Shrock, *Phys. Rev. D* **100**, 073011 (2019).
- [79] U. Mosel, *J. Phys. G* **46**, 113001 (2019).
- [80] J. M. Campbell *et al.*, [arXiv:2203.11110](https://arxiv.org/abs/2203.11110).
- [81] J. Isaacson, S. Höche, D. L. Gutierrez, and N. Rocco, *Phys. Rev. D* **105**, 096006 (2022).
- [82] A. M. Abdullahi *et al.*, [arXiv:2207.04137](https://arxiv.org/abs/2207.04137).
- [83] C. Andreopoulos *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **614**, 87 (2010).
- [84] C. Andreopoulos *et al.*, [arXiv:1510.05494](https://arxiv.org/abs/1510.05494).
- [85] R. Brun and F. Rademakers, *Nucl. Instrum. Methods Phys. Res., Sect. A* **389**, 81 (1997).
- [86] ROOT version 6.18/02.
- [87] L. Alvarez-Ruso *et al.*, *Eur. Phys. J. Special Topics* **230**, 4449 (2021).
- [88] GENIE DarkNeutrino module, accessed 05/10/2022.
- [89] E. Bertuzzo, S. Jana, P. A. N. Machado, and R. Z. Funchal, *Phys. Rev. Lett.* **121**, 241801 (2018).
- [90] M. Drewes, *Int. J. Mod. Phys. E* **22**, 1330019 (2013).
- [91] S. Antusch *et al.*, *J. High Energy Phys.* 10 (2006) 084.
- [92] N. R. Agostinho, G. C. Branco, P. M. F. Pereira, M. N. Rebelo, and J. I. Silva-Marcos, *Eur. Phys. J. C* **78**, 895 (2018).
- [93] C. Soumya, *Phys. Rev. D* **105**, 015012 (2022).
- [94] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou, *J. High Energy Phys.* 09 (2020) 178.
- [95] S. C. Chuliá, R. Srivastava, and A. Vicente, *J. High Energy Phys.* 03 (2021) 248.
- [96] D. Borah and B. Karmakar, *Phys. Lett. B* **789**, 59 (2019).
- [97] J.-L. Tastet and I. Timiryasov, *J. High Energy Phys.* 04 (2020) 005.
- [98] G. Anamiati, M. Hirsch, and E. Nardi, *J. High Energy Phys.* 10 (2016) 010.
- [99] R. A. Gustafson, R. Plestid, and I. M. Shoemaker, *Phys. Rev. D* **106**, 095037 (2022).
- [100] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, *Phys. Rev. D* **98**, 115015 (2018).
- [101] M. Ovchinnikov, T. Schwetz, and J.-Y. Zhu, [arXiv:2210.13141](https://arxiv.org/abs/2210.13141).
- [102] J. Kopp, P. A. N. Machado, M. Maltoni, and T. Schwetz, *J. High Energy Phys.* 05 (2013) 050.
- [103] P. Coloma, P. A. N. Machado, I. Martinez-Soler, and I. M. Shoemaker, *Phys. Rev. Lett.* **119**, 201804 (2017).
- [104] P. Coloma *et al.*, *Eur. Phys. J. C* **81**, 78 (2020).
- [105] W.-Y. Keung and G. Senjanović, *Phys. Rev. Lett.* **50**, 1427 (1983).
- [106] A. G. Hessler, A. Ibarra, E. Molinaro, and S. Vogl, *Phys. Rev. D* **91**, 115004 (2015).
- [107] D. Alva, T. Han, and R. Ruiz, *J. High Energy Phys.* 02 (2015) 072.
- [108] B. Fuks, J. Neundorff, K. Peters, R. Ruiz, and M. Saimpert, *Phys. Rev. D* **103**, 055005 (2021).
- [109] K. Bondarenko *et al.*, *J. High Energy Phys.* 11 (2018) 032.
- [110] R. E. Shrock, *Phys. Rev. D* **24**, 1232 (1981).
- [111] P. Ballett, T. Boschi, and S. Pascoli, *J. High Energy Phys.* 03 (2020) 111.
- [112] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks, *Comput. Phys. Commun.* **185**, 2250 (2014).
- [113] T. Hahn, *Comput. Phys. Commun.* **140**, 418 (2001).
- [114] R. Mertig, M. Böhm, and A. Denner, *Comput. Phys. Commun.* **64**, 345 (1991).
- [115] V. Shtabovenko, R. Mertig, and F. Orellana, *Comput. Phys. Commun.* **207**, 432 (2016).
- [116] V. Shtabovenko, R. Mertig, and F. Orellana, *Comput. Phys. Commun.* **256**, 107478 (2020).
- [117] T. Young and M. J. Mohlenkamp, Introduction to numerical methods and MATLAB programming for engineers (2021).
- [118] FeynRules model database, <https://feynrules.irmp.ucl.ac.be/>, accessed 29-08-2022.
- [119] O. Buss, T. Gaitanos, K. Gallmeister, H. van Hees, M. Kaskulov, O. Lalakulich, A. B. Larionov, T. Leitner, J. Weil, and U. Mosel, *Phys. Rep.* **512**, 1 (2012).
- [120] T. Golan, J. T. Sobczyk, and J. Żmuda, *Nucl. Phys.* **B229–232**, 499 (2012).

- [121] Y. Hayato and L. Pickering, *Eur. Phys. J. Special Topics* **230**, 4469 (2021).
- [122] R. Ruiz, *Phys. Rev. D* **103**, 015022 (2021).
- [123] C. Degrande, O. Mattelaer, R. Ruiz, and J. Turner, *Phys. Rev. D* **94**, 053002 (2016).
- [124] V.M. Gorkavenko, Y.R. Borysenkova, and M.S. Tsarenkova, *J. Phys. G* **48**, 105001 (2021).
- [125] T. Boschi, NuShock: Bunch of Tools to Study Sensitivity of Heavy Neutrinos (2021).
- [126] B. Batell, J. A. Evans, S. Gori, and M. Rai, *J. High Energy Phys.* **05** (2021) 049.
- [127] J. Strait *et al.*, [arXiv:1601.05823](https://arxiv.org/abs/1601.05823).
- [128] A. A. Abud *et al.*, *Instruments* **5**, 31 (2021).
- [129] L. Aliaga *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **743**, 130 (2014).
- [130] M. Makariev (NA49 Collaboration), *AIP Conf. Proc.* **899**, 203 (2007).
- [131] A. Aduszkiewicz *et al.* (NA61/SHINE Collaboration), *Phys. Rev. D* **100**, 112001 (2019).
- [132] L. Aliaga *et al.* (MINERvA Collaboration), *Phys. Rev. D* **94**, 092005 (2016).
- [133] R. Hatcher, Proposal for a Unified “Flux” N-Tuple Format (2012).
- [134] C. Argüelles, P. Coloma, P. Hernández, and V. Muñoz, *J. High Energy Phys.* **02** (2020) 190.
- [135] A. Rubbia, Phenomenology of Particle Physics, *Relativistic Formulation and Kinematics* (Cambridge University Press, Cambridge, England, 2022), pp. 144–183.
- [136] DUNE neutrino flux files generated with G4LBNF, <https://glaucus.crc.nd.edu/DUNEFluxes/>, accessed 23-11-2022.
- [137] A. B. Balantekin, A. de Gouvêa, and B. Kayser, *Phys. Lett. B* **789**, 488 (2019).
- [138] J.-M. Levy, [arXiv:1805.06419](https://arxiv.org/abs/1805.06419).
- [139] A. de Gouvêa, P. J. Fox, B. J. Kayser, and K. J. Kelly, *Phys. Rev. D* **104**, 015038 (2021).
- [140] E. J. Chun, A. Das, S. Mandal, M. Mitra, and N. Sinha, *Phys. Rev. D* **100**, 095022 (2019).
- [141] R. E. Shrock, *Phys. Rev. Lett.* **40**, 1688 (1978).
- [142] E. Gallas *et al.* (FMMF Collaboration), *Phys. Rev. D* **52**, 6 (1995).
- [143] A. Das, P. Konar, and S. Majhi, *J. High Energy Phys.* **06** (2016) 019.
- [144] L. Fischer (IceCube Collaboration), *Proc. Sci. ICHEP2022* (2022) 190.
- [145] G. Bellini *et al.* (Borexino Collaboration), *Phys. Rev. D* **88**, 072010 (2013).
- [146] C. Hagner, M. Altmann, F.V. Feilitzsch, L. Oberauer, Y. Declais, and E. Kajfasz, *Phys. Rev. D* **52**, 1343 (1995).
- [147] M. Verstraeten, Search for sterile neutrinos in the eV and MeV mass range with the SoLid detector, Ph.D. thesis, Antwerp University, 2021.
- [148] A. de Gouvêa and A. Kobach, *Phys. Rev. D* **93**, 033005 (2016).
- [149] E. Cortina Gil *et al.* (NA62 Collaboration), *Phys. Lett. B* **807**, 135599 (2020).
- [150] E. Cortina Gil *et al.* (NA62 Collaboration), *Phys. Lett. B* **816**, 136259 (2021).
- [151] K. Bondarenko, A. Boyarsky, T. Bringmann, M. Hufnagel, K. Schmidt-Hoberg, and A. Sokolenko, *J. High Energy Phys.* **03** (2020) 118.
- [152] A. Boyarsky *et al.*, *Phys. Rev. D* **104**, 023517 (2021).
- [153] D. Boyanovsky, *Phys. Rev. D* **90**, 105024 (2014).
- [154] J.-L. Tastet, O. Ruchayskiy, and I. Timiryasov, *J. High Energy Phys.* **12** (2021) 182.
- [155] K. Bondarenko, A. Boyarsky, J. Klaric, O. Mikulenko, O. Ruchayskiy, V. Syvolap, and I. Timiryasov, *J. High Energy Phys.* **07** (2021) 193.
- [156] M. Drewes, J. Hajer, J. Klaric, and G. Lanfranchi, *J. High Energy Phys.* **07** (2018) 105.
- [157] GENIE BeamHNL module (2023), (Expected official release: 24/Feb/2023).
- [158] GENIE event generator & global analysis of neutrino, <http://www.genie-mc.org/>, accessed 28-08-2022.
- [159] W. H. Press *et al.*, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, Cambridge, 2007)