

Dipole-coupled heavy-neutral-lepton explanations of the MiniBooNE excess including constraints from MINERvA data

N. W. Kamp¹, M. Hostert^{2,3}, A. Schneider¹, S. Vergani⁴,
C. A. Argüelles⁵, J. M. Conrad¹, M. H. Shaevitz⁶, and M. A. Uchida⁴

¹*Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

²*University of Minnesota, Minneapolis, Minnesota 55455, USA*

³*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2J 2W9, Canada*

⁴*University of Cambridge, Cambridge CB3 0HE, United Kingdom*

⁵*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

⁶*Department of Physics, Columbia University, New York, New York 10027, USA*



(Received 9 July 2022; accepted 15 February 2023; published 9 March 2023)

We revisit models of heavy neutral leptons (neutrissimos) with transition magnetic moments as explanations of the 4.8σ excess of electronlike events at MiniBooNE. We first reexamine the preferred regions in the model parameter space to explain MiniBooNE, considering also potential contributions from oscillations due to an eV-scale sterile neutrino. We then derive constraints on the model using neutrino-electron elastic scattering data from MINERvA. To carry out these analyses, we have developed a detailed Monte Carlo simulation of neutrissimo interactions within the MiniBooNE and MINERvA detectors using the `LeptonInjector` framework. This simulation allows for a significantly more robust evaluation of the neutrissimo model compared to previous studies in the literature—a necessary step in order to begin making definitive statements about beyond the Standard Model explanations of the MiniBooNE excess. We find that MINERvA rules out a large region of parameter space, but allowed solutions exist at the 2σ confidence level. A dedicated MINERvA analysis would likely be able to probe the entire region of preference of MiniBooNE in this model.

DOI: [10.1103/PhysRevD.107.055009](https://doi.org/10.1103/PhysRevD.107.055009)

I. INTRODUCTION

Past, current, and future neutrino experiments offer some of the most promising avenues for observing physics beyond the Standard Model (BSM). The discovery of neutrino oscillations, and thus nonzero neutrino mass, is itself an indication of BSM physics [1,2], spurring a decades-long oscillation experimental program spanning orders of magnitude in energy and length scales [3–5]. A nearly-consistent three-neutrino mixing paradigm has emerged from this program; however, anomalous results have also been observed. Two striking examples are the excess of inverse-beta-decay events at the LSND detector at the Los Alamos Neutrino Science Center [6] and the excess of electronlike events observed by the MiniBooNE (MB) experiment at the Fermilab Booster Neutrino Beam (BNB) [7]. Determining the nature of these excesses is an active frontier in neutrino physics [8]. Historically, both

excesses have been interpreted within the context of a $3 + 1$ model, in which one introduces an eV-scale sterile neutrino facilitating short-baseline $\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_e(\bar{\nu}_e)$ oscillations. However, models beyond the vanilla $3 + 1$ scenario may better accommodate these anomalies within the global experimental landscape, including cosmology and other neutrino oscillation experiments.

One such model considers MeV-scale heavy neutral leptons (HNLs), or “neutrissimos” [9,10], with transition magnetic moments, popularly known as the dipole portal to HNLs [11–24]. Most recently, Ref. [23] showed that the MB anomaly could be fitted with such a dipole portal extended to include a mass-mixed eV-scale sterile neutrino. The advantages of this model are twofold: (1) it provides a better description of the low-energy, forward-angle part of the MB excess compared to the $3 + 1$ model alone, and (2) reducing the oscillation-based contribution to the MB excess alleviates tension in global fits to the $3 + 1$ model while retaining an explanation to the LSND anomaly.

For the main result of this paper, we derive constraints on the dipole portal coupling using existing experimental results from the MINERvA Collaboration. Specifically, MINERvA has performed measurements of the neutrino-electron elastic scattering ($\nu - e$) rate for the purpose of

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. 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.

constraining the Neutrino Main Injector (NuMI) low-energy (LE) and medium-energy (ME) neutrino fluxes [25–27]. Photons produced in the decay of the dipole-coupled neutrino would mimic the single electromagnetic shower morphology of electron-scattering (ES) events and would therefore enter as a photonlike (large dE/dx) background in the MINERvA analysis. In addition to deriving new constraints from MINERvA, we improve upon the MiniBooNE dipole model analysis performed in Ref. [23].

The analyses described in this article have been carried out using a novel simulation developed within the `LeptonInjector` framework [28]. This tool allows a more robust description of neutrissimo interactions in detector subsystems as well as the position and kinematics of observable final state particles. This is an important improvement over previous treatments in the literature, as an accurate simulation of exotic BSM physics scenarios in different neutrino detectors will be vital in determining the nature of the MiniBooNE excess.

The rest of this article is organized as follows. In Sec. II, we review the dipole-portal sterile neutrino model in more detail. In Sec. III, we introduce the novel `LeptonInjector`-based simulation developed for the studies presented here. In Sec. IV, we refine the preferred regions in dipole parameter space which explain the energy and angular distributions of the MiniBooNE excess. In Sec. V, we calculate constraints in dipole parameter space derived from the MINERvA ES analysis. In Sec. VI, we discuss existing and projected constraints on the dipole model from current and future neutrino experiments. Finally, in Sec. VII we discuss the outlook of this mixed model of oscillation and decay as a solution to the MiniBooNE anomaly in light of our derived MINERvA constraints.

II. THE DIPOLE PORTAL

We introduce a right-handed neutrino \mathcal{N} that couples to the left-handed neutrino fields via a transition magnetic moment. At the effective operator level, we have the dimension-six dipole operators,

$$\mathcal{L} \supset \frac{1}{\Lambda^2} \bar{L}_\alpha \tilde{H} \sigma^{\mu\nu} \mathcal{N}_R (C_B^\alpha B_{\mu\nu} + C_W^\alpha W_{\mu\nu}^a \sigma_a) + \text{H.c.}, \quad (1)$$

where α is a flavor index, $\tilde{H} = i\sigma_2 H^*$, C_B^α , and C_W^α are Wilson coefficients and Λ is the new physics scale. After electroweak symmetry breaking, the dipole operator gives rise to the electromagnetic transition magnetic moment of neutrinos,

$$\mathcal{L} \supset d_{\alpha\mathcal{N}} \bar{\nu}_\alpha \sigma_{\mu\nu} F^{\mu\nu} \mathcal{N}_R + \text{H.c.}, \quad (2)$$

where ν_α corresponds to the neutrino Weak eigenstates and $F^{\mu\nu}$ to the electromagnetic field strength. The dipole parameter is defined as $d_{\alpha\mathcal{N}} = (v_h/\sqrt{2})(C_W^\alpha C_B^\alpha + C_{W,S}^\alpha)/\Lambda^2$, where v_h is the vacuum expectation value of

the SM Higgs. The other transition moments mediated by the W and Z bosons will also be present, but their low-energy effects are further suppressed by G_F , and therefore negligible in our region of interest.

The upscattering signature we are interested in at MiniBooNE is initiated by muon neutrinos and antineutrinos, so in this work, we always consider the second-generation coupling $d_{\mu\mathcal{N}}$. In UV completions of this operator from one-loop diagrams, the size of the dipole coupling is typically proportional to $d_{\alpha\mathcal{N}} \sim (1/16\pi^2)(m_\beta/m_X^2)$, where m_β is the mass of some charged particle and X is some heavy, charged scalar, for example. Under the assumption of flavor-conserving interactions between SM neutrinos and the new physics, one would take $\beta = \alpha$ and conclude that the transition magnetic moment of the third-generation neutrinos is much larger. While this need not necessarily be the case, it still provides enough motivation for us to consider the third-generation coupling $d_{\tau\mathcal{N}}$, including the case

$$\frac{d_{\tau\mathcal{N}}}{d_{\mu\mathcal{N}}} = \frac{m_\tau}{m_\mu}, \quad (3)$$

and neglecting the first-generation couplings altogether.

We note that large neutrino magnetic moments typically imply large Dirac masses for light neutrinos, $m_D \bar{\nu}_L \mathcal{N}_R$ [29–31]. In models with heavy neutrinos, this presents two challenges: (i) neutrino masses, schematically given by $m_\nu \sim m_D^2/M$, may be too large, and (ii) m_D will generate large mixing between active and heavy neutrinos, $U_{\alpha\mathcal{N}} \sim m_D/M$. Here, M stands for the heavy neutrino mass scale. The first challenge is easily overcome in models like the inverse-seesaw where lepton number is approximately conserved and right-handed neutrino fields combine into pseudo-Dirac \mathcal{N} particles. In that case, m_ν is controlled by the mass splitting between the Majorana neutrinos, which may be parametrically small. This is also the case preferred by the short-baseline phenomenology discussed below, since the single-photons produced in the decays of Dirac HNLs are less forward. The second challenge, however, is not so easy to overcome. The mixing between active and heavy neutrinos remains large even in inverse-seesaw models and the parameter space in that case is strongly constrained by laboratory limits on $U_{\alpha\mathcal{N}}$. For instance, decay-in-flight signatures at neutrino experiments, where \mathcal{N} is copiously produced in meson decays at the target, can set limits as strong as $|U_{\alpha\mathcal{N}}|^2 < \mathcal{O}(10^{-11})$ [32].

There are several models in the literature that can suppress the Dirac mass in comparison with the magnetic moment of neutrinos [33–39], but only to a certain extent and not without fine tuning. We proceed assuming that m_D is sufficiently small so as not to impact the phenomenology, but note that depending on the amount of fine tuning, bounds on the mixing angles $|U_{\alpha\mathcal{N}}|$ would also need to be considered.

In performing a fit to MiniBooNE data, we also include a sterile neutrino, ν_s , which mixes with light neutrinos but does not have transition magnetic moments with active neutrinos. This sterile will be responsible for short-baseline oscillations with Δm_{41}^2 of $\mathcal{O}(1 \text{ eV}^2)$. We provide two examples, the case of a global fit to short-baseline data excluding MiniBooNE [23] and the case of a joint fit to only MiniBooNE and the recent MicroBooNE CCQE-like analysis [40]. It is also important that ν_s and the heaviest neutrino do not mix, as otherwise \mathcal{N}_R would also mix with light neutrinos via ν_s . In summary, our spectrum is defined as

$$\nu_i \sim \sum_{\alpha=\{e,\mu,\tau,s\}} U_{\alpha i} \nu_\alpha, \quad \text{for } i \leq 4 \quad \text{and} \quad \nu_5 \equiv \mathcal{N}, \quad (4)$$

where the approximate symbol means that any additional term contains very small mixing elements. In addition, this mixing would result in a corresponding dipole coupling $d_{\alpha s} \bar{\nu}_\alpha \sigma_{\mu\nu} F^{\mu\nu} \nu_s$, which is strongly constrained by big bang nucleosynthesis and stellar cooling. While the former may be modified *à la* secret interactions, the latter is significantly more robust and constrains $d_{\mu s} < 2 \times 10^{-12} \mu_B$ [41].

The relevant interactions for this work are shown in Fig. 1, which include Primakoff upscattering off of a nuclear target $\nu A \rightarrow \mathcal{N} A$ (left) and the radiative decay of the heavier neutrino $\mathcal{N} \rightarrow \nu_\alpha \gamma$ (right). Due to the photon propagator, the scattering process is dominated by coherent exchange with the nucleus, except at the largest $m_{\mathcal{N}}$ values. We include helicity-flipping upscattering, where the helicity of \mathcal{N} is opposite to that of ν_α , and helicity-conserving, where they are the same. The latter is suppressed by the typical energy of the process, $m_{\mathcal{N}}^2/E^2$, and is only relevant at the largest masses. To give an example, the cross section for an incoming neutrino with energy $E_{\nu_\alpha} = 1 \text{ GeV}$ upscattering off of a carbon nucleus into a neutriissimo with mass $m_{\mathcal{N}} = 50 \text{ MeV}$ and dipole coupling $d_{\alpha \mathcal{N}} = 10^{-6} \text{ GeV}^{-1}$ is $\sigma \approx 5 \times 10^{-39} \text{ cm}^2$ [18].

The rest frame decay width of the \mathcal{N} is given by the incoherent sum over all outgoing flavors,

$$\Gamma_{\mathcal{N} \rightarrow \nu \gamma} = \left(\sum_{\alpha} |d_{\alpha \mathcal{N}}|^2 \right) \frac{m_{\mathcal{N}}^3}{4\pi}. \quad (5)$$

Throughout this work, we will assume that either \mathcal{N} or light neutrinos are Dirac particles, such that the differential decay




FIG. 1. \mathcal{N} production from ν upscattering (left) and \mathcal{N} decay (right).

rate is proportional to $(1 \pm \cos \theta)$ [19]. For the region of dipole parameter space preferred by MiniBooNE [23], the HNL has a lab frame decay length $L \sim 1\text{--}10 \text{ m}$ for typical MINERvA neutrino energies and $L \leq 1 \text{ m}$ for typical MiniBooNE neutrino energies. Thus, Primakoff upscattering near or within each detector is the most relevant production mechanism in this region of parameter space. For smaller HNL masses, decay lengths become longer and upscattering in the dirt along the BNB and NuMI beamlines becomes more important for MiniBooNE and MINERvA, respectively.

III. SIMULATING NEUTRISSIMOS

In order to describe neutriissimo interactions in the MiniBooNE and MINERvA detectors, we have developed a custom simulation based in the `LeptonInjector` framework [28]. The simulation begins by injecting neutrinos in MiniBooNE and MINERvA according the BNB and NuMI beam profiles, respectively. For each neutrino, a flight path is randomly selected within a cone surrounding the detector. This path is then used to calculate intersections with different components of the beamline, including bedrock between the target and detector as well as various detector subsystems. An upscattering location is sampled along this flight path according to the cross section in each traversed material. The final state kinematics of the produced HNL are sampled according to the $d\sigma/dy$ distribution, where $y \equiv (E_\nu - E_{\mathcal{N}})/E_\nu$ in the lab frame [41]. See the Appendix for more details on the upscattering cross section used in this analysis. Next, a decay location is sampled along the flight path of the HNL. The final state kinematics of decay photon are sampled according to $d\Gamma/d\cos\theta \propto (1 - \cos\theta)$. We trace the flight path of the photon through different materials until it converts to an e^+e^- pair, which are assumed to reconstruct as a single electromagnetic shower within the MiniBooNE and MINERvA detectors.

The above procedure is easily generalizable to other BSM scenarios and detector configurations. The user needs only to provide (1) the flux of initial state particles (neutrinos in this case), (2) the total and differential cross sections for the relevant processes, and (3) the relevant detector geometry. The flux and cross section can be input either as analytic expressions or splines; `LeptonInjector` will interpolate in the latter case. The detector geometry is set using a simple configuration file in which the user instantiates any number of volumes. `LeptonInjector` supports most simple volumes as well as general extruded polygons—for example, the nuclear targets in MINERvA shown in Fig. 2, which are subsections of hexagonal prisms. `LeptonInjector` will then generate user-specified initial state particles, sample their interaction locations, and store the kinematics and weights of the final state particles. The simulation is set up to be as efficient as possible, such that most generated events create an observable final state within

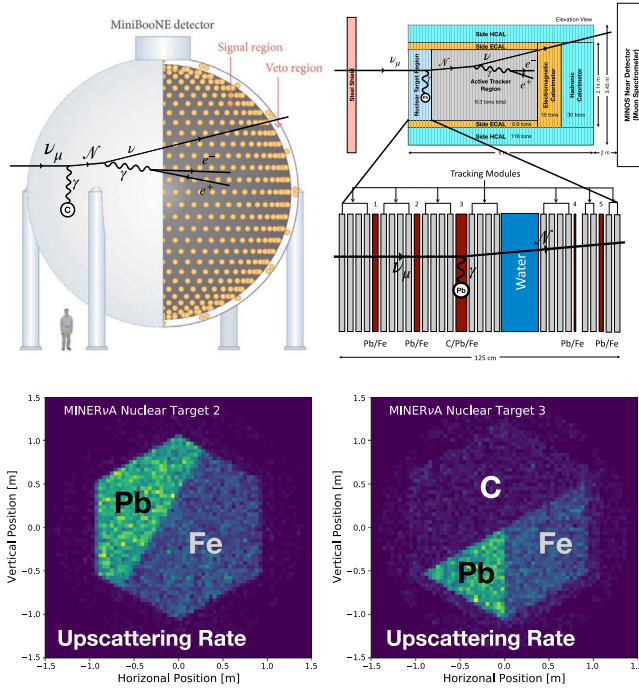


FIG. 2. (Top) Schematic representation of HNL production via upscattering and the subsequent HNL decay within the MiniBooNE (left) and MINERvA (right) detectors. Detector images have been adapted from Refs. [46,47]. (Bottom) Example upscattering rates within two of the MINERvA nuclear targets as simulated using *LeptonInjector*. The coherent enhancement of the upscattering cross sections leads to a larger rate in the high- Z components.

the fiducial volume and are appropriately downweighted. In the analysis presented here, we have tracked final state particles until the production of an observable e^+e^- pair. This strategy can be adapted for other BSM scenarios which involve multiple interactions before the production of an observable final state.

The simulation performed for the analysis presented here gives a robust estimation of the e^+e^- event rate within the fiducial volume of each experiment. Further, it provides an accurate kinematic description of the photons which survive fiducialization (namely, the photon energy and angle with respect to the beamline). This is vital, as it will be shown that the region in dipole parameter space preferred by the MiniBooNE excess is highly sensitive to the kinematics of the final state photon. Such an effect has been appreciated by previous studies in the literature [11,24], but a comprehensive fit over the full dipole parameter space was not attempted until Ref. [23]; the simulation described above allows us to refine the fit from Ref. [23].

The simulation is also vital in order to properly assess the ability of MINERvA to constrain the dipole solution to the MiniBooNE excess. By leveraging our simulation, we will show that the reconstruction efficiency for radiative HNL decays in MINERvA varies by 3–4 orders of magnitude

across the parameter space. Previous calculations in the literature have assumed a constant 10% reconstruction efficiency for radiative HNL decays in MINERvA [42], and have thus incorrectly concluded that MINERvA data rules out the dipole-portal MiniBooNE solution. Further, they do not simulate the complex subcomponents of the MINERvA detector—this is an important step, as upscattering in the high- Z nuclear targets just upstream of the fiducial volume can contribute significantly to the single shower sample for short-lived HNLs.

Thus, the *LeptonInjector*-based simulation performed for this study is essential to evaluate the status of the dipole-coupled neutrissimos as an explanation for the MiniBooNE excess. As the community turns toward more exotic BSM explanations of the MiniBooNE excess, it is imperative that these models are evaluated within the context of realistic detector descriptions. Our simulation framework is the perfect tool for such a task; though the version used in this article was developed specifically to study neutrino interactions in MiniBooNE and MINERvA, it can be adapted easily to accommodate additional BSM scenarios and neutrino detectors.

IV. NEUTRISSIMOS AT MiniBooNE

The MiniBooNE detector uses Cherenkov light to detect final state particles produced in neutrino interactions. As electrons and photons both produce electromagnetic showers which show up as distorted Cherenkov rings in the detector, the two particles are indistinguishable in MiniBooNE. Thus, photons from the dipole model could contribute to the MiniBooNE electronlike excess. In the case of a nonzero effective dipole coupling, $d_{\mu N}$, muon neutrinos from Fermilab’s Booster Neutrino Beam will undergo Primakoff upscattering into HNL states. This can happen via photon exchange with nuclear targets both within the dirt between the BNB target and the detector and within the CH_2 detector volume itself.

A. Simulation

As described in Sec. III, we use *LeptonInjector* to simulate the production and decay of HNLs in MiniBooNE [28]. A schematic depiction of this process is shown in the top panel of Fig. 2. Muon neutrinos are injected according to the BNB flux and allowed to upscatter to HNLs along the 541 m baseline between the BNB target and the detector or within the 818-ton CH_2 detector itself. We describe the MiniBooNE detector as a sphere of CH_2 with a total radius of 6.1 m and a fiducial radius of 5 m. The MiniBooNE detector sits within a sphere of air with a radius of 9 m, meant to represent the detector hall. To simulate the BNB, muon neutrinos generated in *LeptonInjector* propagate through 541 m of dirt before reaching the MiniBooNE detector. The neutrino upscattering and subsequent HNL decay are

simulated according to the procedure outlined in Sec. III. We parametrize the postfiducialization photon detection efficiency as a linearly decreasing function of the true photon kinetic energy [23]. We also impose a reconstruction threshold on the photon kinetic energy of 140 MeV [43]. We also independently smear the reconstructed visible energy and scattering angle of each photon according to the resolution of each as a function of true photon kinetic energy. The energy (angular) resolution comes from a power-law (quadratic) fit to simulated single electromagnetic shower events in MiniBooNE [44], with typical values of 10% (3°) for photons from this model, consistent with figures reported by the MiniBooNE Collaboration [45].

B. Analysis methodology

The dipole model has previously been studied as a potential solution to the MiniBooNE anomaly [11–18,22–24]. Specifically, we expand upon the study performed in Ref. [23], which examined a mixed model consisting of an eV-scale ν_4 facilitating short-baseline $\nu_\mu \rightarrow \nu_e$ oscillations as well as an MeV-scale HNL decaying to a photon via the dipole portal mechanism. Ref. [23] found a preference for a dipole-coupled HNL with $m_{\mathcal{N}} \sim 400$ MeV and $d_{\mu\mathcal{N}} \sim 3 \times 10^{-7}$ GeV $^{-1}$.

We make a number of improvements to this analysis. First, we make use of a more robust nuclear electromagnetic form factor, implementing a data-driven Fourier-Bessel function parametrization [48–50] with the data files made available in Ref. [51]. This should be compared with the simpler dipole parametrization, which overestimates the differential cross sections at larger momentum exchange (see the Appendix for more details). This has an impact on the allowed regions at larger heavy neutrino masses, as the form factor used in this study drops off much more quickly at larger Q^2 . It also reduces the contribution from the dipole model at large scattering angles, thereby making it difficult to explain the backscattered lepton angular distribution of the MiniBooNE excess. This effect has been pointed out in previous studies of the dipole model in MiniBooNE [16]. In this work, we perform a more detailed analysis of the dipole parameter space to determine whether solutions exist which can accommodate both the energy and angular distributions of the excess. The statistical treatment for each distribution has been improved—we now consider correlated systematic errors in the reconstructed E_ν^{QE} distribution from the provided covariance matrix (after constraining with the covariance matrix for MiniBooNE’s ν_μ dataset). We also introduce an uncorrelated systematic error of 13% in the $\cos\theta$ distribution, consistent with that in the E_ν^{QE} distribution.

We perform fits to the excess only in neutrino-mode data, as MiniBooNE has collected about an order of magnitude more events in this beam configuration compared to their antineutrino-mode data. We use the simulated photon events

from the above procedure to perform two different spectral analyses across dipole parameter space: one in the E_ν^{QE} distribution and one in the $\cos\theta$ distribution. In both cases, we calculate a χ^2 test statistic comparing the dipole model prediction to the remaining excess after subtracting off the oscillation contribution from the MiniBooNE-less global fit reported in Ref. [23]. In the E_ν^{QE} fit, we use the electronlike channel fractional covariance matrix provided by the MiniBooNE Collaboration after constraining with the covariance matrix in the muonlike channel. No systematic errors are provided for the $\cos\theta$ distribution; therefore, as mentioned above, we consider an uncorrelated fractional systematic error of 13% in each bin of the $\cos\theta$ prediction, consistent with the level in the E_ν^{QE} channel. Confidence regions are drawn using a $\Delta\chi^2$ test statistic, assuming Wilks’ theorem with two degrees of freedom [52].

C. Results

The result from the fit procedure described above is shown in Fig. 3. One can see that it is difficult to explain the E_ν^{QE} and $\cos\theta$ distributions through the same dipole-coupled HNL, as the two distributions prefer different regions of dipole parameter space. The preferred regions overlap at the 2σ C.L., though some of this overlap region is in tension with constraints derived from the NOMAD single-photon analysis [53]. As stated above, the difference between this result and the result in Ref. [23] is driven mainly by the updated form factor. This reduces the Primakoff upscattering rate at large scattering angles,




FIG. 3. The 2σ and 3σ C.L. preferred regions to explain the MiniBooNE anomaly in mass-coupling parameter space for a dipole-coupled heavy neutral lepton. The pink (green) curves correspond to results from fitting the E_ν^{QE} ($\cos\theta$) distribution. Dipole model fits are performed after subtracting the oscillation component from a global fit to a $3+1$ model excluding MiniBooNE data [23]. In general, the energy and angular distributions prefer different regions of parameter space, though overlap exists at the $2\sigma - 3\sigma$ level.

requiring the fit to push to larger \mathcal{N} masses and dipole couplings to explain this part of the MiniBooNE excess. That being said, there is a region in parameter space for $d \sim 10^{-6}$ GeV $^{-1}$ and $m_{\mathcal{N}} \sim 0.5$ GeV which is (1) within the 2σ allowed region for the E_{ν}^{QE} distribution, (2) within the 2σ allowed region for the $\cos\theta$ distribution, and (3) not ruled out by NOMAD's single-photon search [53]. In Fig. 4 we show E_{ν}^{QE} and $\cos\theta$ distributions for an example parameter point in this region, located at $d = 1.25 \times 10^{-6}$ GeV $^{-1}$ and $m_{\mathcal{N}} = 0.47$ GeV. One can see that this model can describe most of the excess except for the region at $\cos\theta \lesssim 0$.

This situation might be improved when considering alternative oscillation scenarios. The above fits assumed the MiniBooNE-less global-fit result, which found a best-fit solution at $\{\Delta m^2 \approx 1.3$ eV 2 , $\sin^2(2\theta_{\mu e}) \approx 6.9 \times 10^{-4}\}$ [23]. We now consider an alternative sterile neutrino hypothesis:

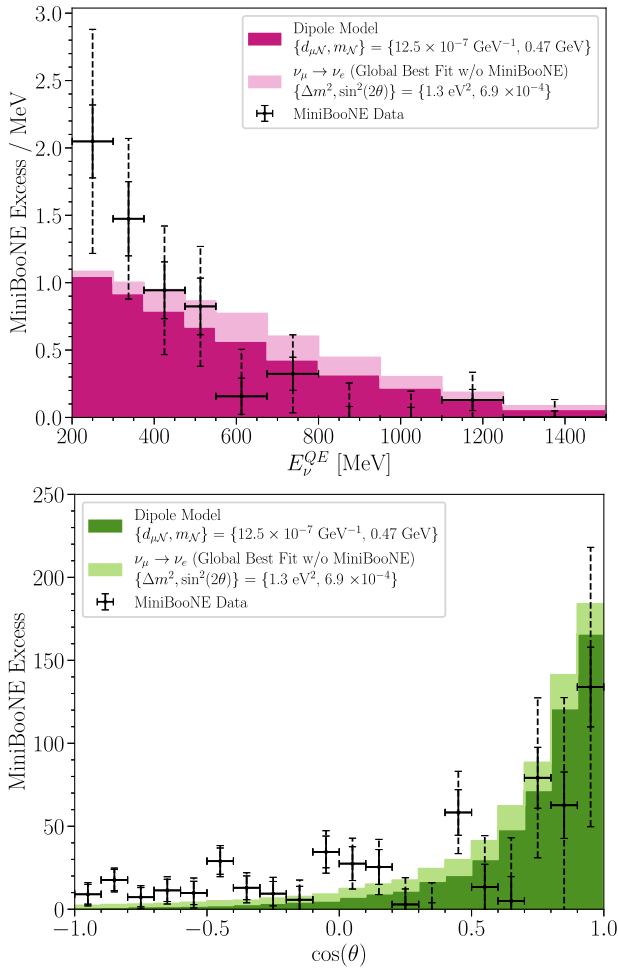


FIG. 4. The E_{ν}^{QE} (top) and $\cos\theta$ distributions at the example dipole model hypothesis indicated by the black star in Fig. 3. The darker contribution in each stacked histogram corresponds to the dipole model prediction, while the lighter contribution corresponds to the oscillation contribution. The background-subtracted MiniBooNE excess is indicated by the black data points, with solid and dashed error bars indicating statistical and statistical + systematic errors, respectively.

the result from the recent MiniBooNE + MicroBooNE CCQE-like 3 + 1 combined fit performed by the MiniBooNE Collaboration [40].

The MiniBooNE + MicroBooNE CCQE-like 3 + 1 combined analysis found a best-fit solution at $\{\Delta m^2 \approx 0.2$ eV 2 , $\sin^2(2\theta_{\mu e}) \approx 0.03\}$ [40]. This introduces a much larger $\nu_{\mu} \rightarrow \nu_e$ oscillation component in MiniBooNE. Thus, the dipole model is primarily driven to explain the lowest energy and most forward-angle portion of the excess. A mild preference for a dipole-coupled heavy neutral lepton is found at the 1σ level. As the dipole model is no longer required to explain the broad-angle portion of the MiniBooNE excess, the angular fit is able to accommodate a large range of heavy neutrino masses while the energy fit prefers lower heavy neutrino masses at $m_{\mathcal{N}} \lesssim 100$ MeV. The preferred regions in dipole parameter space under this oscillation hypothesis are shown in Fig. 5. As a benchmark point, we consider a solution at $d = 1.7 \times 10^{-7}$ GeV $^{-1}$ and $m_{\mathcal{N}} = 0.08$ GeV. As shown in Fig. 6, this benchmark point can reasonably describe the E_{ν}^{QE} and $\cos\theta$ distributions of the MiniBooNE excess.

One can also consider a nonzero transition magnetic moment coupling between the \mathcal{N} and the ν_{τ} flavor eigenstate. This would open up the decay channel $\mathcal{N} \rightarrow \nu_{\tau}\gamma$, increasing the decay width by the ratio $(|d_{\mu\mathcal{N}}|^2 + |d_{\tau\mathcal{N}}|^2)/|d_{\mu\mathcal{N}}|^2$. This will have a more pronounced impact on the fit in the lower HNL mass region of parameter space, as lifetimes in the higher HNL mass region are sufficiently




FIG. 5. 1σ C.L. preferred and 2σ C.L. allowed regions with regard to the MiniBooNE anomaly in mass-coupling parameter space for a dipole-coupled heavy neutral lepton. The pink (green) curves correspond to results from fitting the E_{ν}^{QE} ($\cos\theta$) distribution. Dipole model fits are performed after subtracting the oscillation component from the combined MiniBooNE + MicroBooNE CCQE-like 3 + 1 fit [40]. Mild preference for a dipole-coupled heavy neutrino with $m_{\mathcal{N}} \lesssim 100$ MeV is found at the 1σ level.




FIG. 6. Similar to Fig. 4. The E_{ν}^{QE} (top) and $\cos\theta$ (bottom) distributions at the example dipole model hypothesis indicated by the black star in Fig. 5. The oscillation component comes from the combined MiniBooNE + MicroBooNE CCQE-like 3 + 1 fit [40].

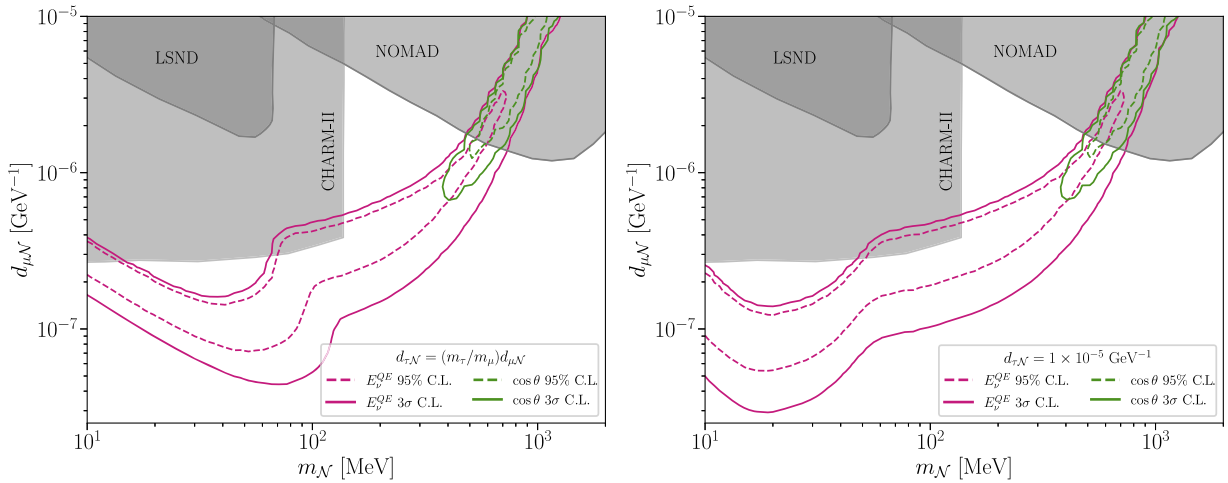


FIG. 7. Similar to Fig. 3, but considering the $\mathcal{N} \rightarrow \nu_{\tau}\gamma$ decay channel. On the left, we show the case with coupling $d_{\tau\mathcal{N}} = (m_{\tau}/m_{\mu})d_{\mu\mathcal{N}}$ and on the right we consider $d_{\tau\mathcal{N}} = 1 \times 10^{-5} \text{ GeV}^{-1}$ instead. Compared to Fig. 3, preferred regions in the lower HNL mass region move to lower $d_{\mu\mathcal{N}}$ values.

short such that introducing another decay channel does not appreciably change the phenomenology.

As discussed in Sec. II, in some UV completions of the dipole model, a natural scaling given by $d_{\tau\mathcal{N}}/d_{\mu\mathcal{N}} = m_{\tau}/m_{\mu}$. The resulting preferred regions in dipole model parameter space under this assumption are shown on the left panel of Fig. 7. One can see that, compared with Fig. 3, solutions explaining the E_{ν}^{QE} distribution have opened at lower $d_{\mu\mathcal{N}}$ couplings for $m_{\mathcal{N}} \lesssim 100 \text{ MeV}$. We also examine the effect of large tau coupling $d_{\tau\mathcal{N}} = 1 \times 10^{-5} \text{ GeV}^{-1}$, which is meant to capture the extent of flexibility introduced into the dipole model when allowing for nonzero $d_{\tau\mathcal{N}}$. The preferred regions for this case are shown in the right panel of Fig. 7. For both cases, we consider an oscillation contribution given by the MiniBooNE-less global fit.

V. NEUTRINO SCATTERING AT MINERvA

Neutrino upscattering can also occur in the MINERvA detector. We choose to study MINERvA for two main reasons: (i) the NuMI beam is a higher-energy beam in comparison with the BNB. This is specially true for the medium-energy (ME) NuMI configuration, where $\langle E_{\nu} \rangle \simeq 7 \text{ GeV}$, but it is still the case for the low-energy (LE) configuration, where $\langle E_{\nu} \rangle \simeq 3 \text{ GeV}$. This allows us to probe HNLs of larger masses. (ii) it is one of the few accelerator experiments in the few GeV region to have a dedicated neutrino-electron ($\nu - e$) scattering analysis. While measurements of this channel have been performed with greater precision at experiments like CHARM [54] and CHARM-II [55], LSND [56], reactors [57], Borexino [58,59], and Super-Kamiokande [60], they are not as well suited for the study of the MiniBooNE explanations considered here, where HNLs have hundreds of MeV in mass. With the exception of CHARM and

CHARM-II, the previous experiments operate at energies below the HNL production threshold and are therefore not sensitive to our region of interest. While we could also consider CHARM and CHARM-II, we note that they observe larger neutrino-induced backgrounds thanks to the faster growth of the SM cross section with respect to the dipole one. In addition, the HNLs would be produced with a larger boost factor, and therefore escape more often.

In this work, we will consider three existing measurements of the neutrino-electron cross section by MINERvA [25–27]. The first was performed in the LE configuration of the NuMI beam operating with a forward-horn current (FHC), optimizing the number of neutrinos produced. The last two were performed in the ME configuration, one in FHC and the other in reverse-horn current (RHC) mode, the latter optimizing the number of antineutrinos. The ME RHC measurement, also the most recent, is particularly sensitive due to the smaller antineutrino- and neutrino-induced backgrounds. Unlike the dipole cross section, antineutrino-nucleus weak cross sections are smaller than neutrino-nucleus cross sections.

A. Simulation

As described in Sec. III, we use `LeptonInjector` [28] to simulate $\nu A \rightarrow \mathcal{N}A$ upscattering inside as well as outside the MINERvA detector. A schematic depiction of this process is shown in the top panel of Fig. 2. We include upscattering in the upstream dirt, in the surrounding air, in the nuclear target planes (detailed below), in the plastic scintillator, as well as in the outermost electromagnetic calorimeter. While most of these components are not part of the fiducial volume for the $\nu - e$ analysis, they can significantly contribute to the signal rate due to the displaced decays of the HNLs. For long-lived HNLs, upscattering in the dirt dominates the signal rate, followed by the nuclear target planes, which contain high-density materials like ^{208}Pb and ^{56}Fe . Detailed modeling of the detector geometry and material composition is necessary to correctly predict the contributions from these different upscattering sites. We define each nuclear target to be a hexagonal prism with apothem 92 cm, matching that of the tracker region [61], and z extent given by Table 4 of Ref. [47]. The six nuclear target planes are detailed in Table I. In the bottom panel of Fig. 2, we show the positional distribution of the upscattering rate within two of the MINERvA nuclear targets as simulated using `LeptonInjector` [28]. We also consider upscattering within the electromagnetic calorimeter, defined as a hexagonal prism surrounding the inner detector with 107 cm apothem, and within the steel veto shield ~ 1 m in front of MINERvA [61]. The fiducial volume of MINERvA for the $\nu - e$ analyses is assumed to be approximately the same for both the LE and ME analyses and is defined as a hexagon of 81 cm apothem with ~ 2.8 m z extent inside the plastic scintillator.

After production, we track the HNL's path through the detector and force a decay to occur before the end of

TABLE I. Specifications of the MINERvA nuclear targets as implemented in `LeptonInjector` [28] for this analysis. Each nuclear target is defined as a hexagonal prism with an apothem of 92 cm. Z positions and extents of each nuclear target have been taken from Table 4 of Ref. [47]. The coordinate system is defined such that $z = 0$ corresponds to the front of the MINERvA detector. We have confirmed that the fiducial mass (bounded by an 85 cm apothem hexagon) of each nuclear target sub-component matches the fiducial mass quoted in Table 4 of Ref. [47]. The last column of this table refers to the mass of each nuclear target subcomponent within the 92 cm apothem hexagonal prism.

Target	z -location (cm)	z -extent (cm)	Mass (kg)
1-Fe	13.6	2.567	370
1-Pb	13.6	2.578	317
2-Fe	31.3	2.563	370
2-Pb	31.3	2.581	317
3-Fe	53.4	2.573	197
3-Pb	53.4	2.563	141
3-C	53.4	7.620	194
Water	89.5	18.06	530
4-Pb	125.6	0.795	263
5-Fe	138.9	1.289	186
5-Pb	138.9	1.317	162

the fiducial volume; we then downweight the event by the probability of decaying within the considered region. Each decay produces a photon for which we physically sample a pair-production location. Events that do not pair produce within the fiducial volume are removed. This procedure accounts for events where the HNL decays outside the fiducial volume, but the photon conversion happens inside of it. This effect is important for short-lived HNLs since the rate of HNLs produced in the high-density lead planes can significantly contribute to the signal rate even though they are not contained in the fiducial volume.

The neutrino fluxes for the LE mode have been taken from Ref. [62] and for the ME they have been digitized from Ref. [63]. The total exposures used in the three $\nu - e$ analyses are 3.43×10^{20} POT for LE-FHC, 1.16×10^{21} POT for ME-FHC, and 1.22×10^{21} POT for ME-RHC.

To reduce neutrino-induced backgrounds, MINERvA applies an extensive list of selection cuts. To properly estimate the resulting efficiency of these cuts in our HNL signal, it is important to correctly model the reconstruction of the energy and angle of the single photons. In the absence of a full detector simulation, we proceed to approximate the detector energy and angular resolutions as Gaussian functions. For the energy resolution, we take $\sigma_E/E = 5.9\% / (\sqrt{E/\text{GeV}}) + 3.4\%$ [25], while for the angular resolution, we take an energy-independent angular resolution of $\sigma_\theta = 0.7^\circ$, assumed to be isotropic in the shower's

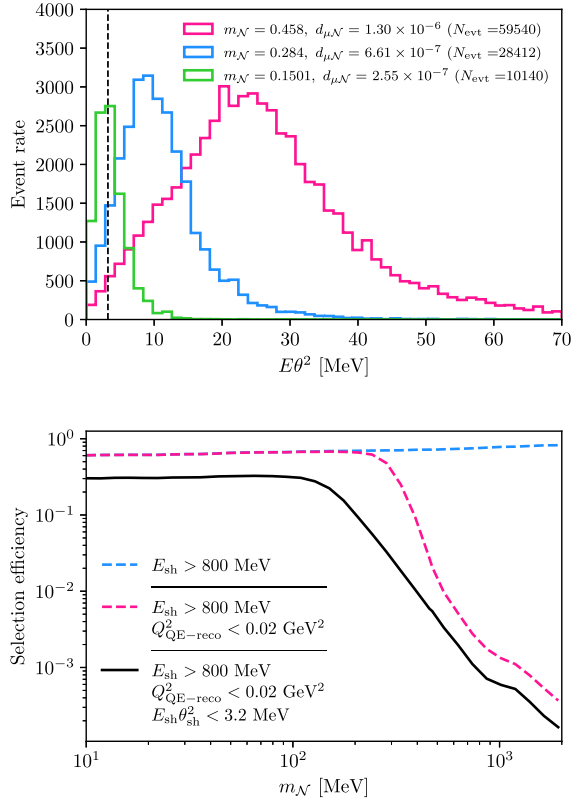


FIG. 8. (Top) The predicted $E_{\text{sh}}\theta_{\text{sh}}^2$ distribution before detector smearing and signal selection for three choices of model parameters at MINERvA. Bottom) The signal selection efficiency of our analysis cuts, excluding the dE/dx cut, as a function of the HNL mass.

azimuthal angle.¹ The angular resolution is implemented by sampling a polar angle $\delta\theta$ from a Gaussian distribution of standard deviation σ_θ , rotating the photon by $\delta\theta$ with respect to its momentum, assigning it an azimuthal angle φ from the uniform distribution $[0, 2\pi]$, and finally rotating the photon back to the laboratory frame by its original polar angle θ_{true} .

B. Event selection

Now we discuss the most important signal selection cuts. The analyses [25,26] make use of a long list of signal selection cuts, designed to suppress as many neutrino-nucleus scattering backgrounds as possible. The most worrisome backgrounds include π^0 production and ν_e CC scattering. The former is particularly important for our radiative decay signal, as it can give rise to coherent single-photonlike signatures. Cuts related to the shower radius and transversal as well as longitudinal profiles are not implemented in our analysis but are expected to have

¹This is only an approximation, as the MINERvA detector is not azimuthally symmetric. Nevertheless, the differences in resolution in the X and Y planes are small [61,64] and neglected here.

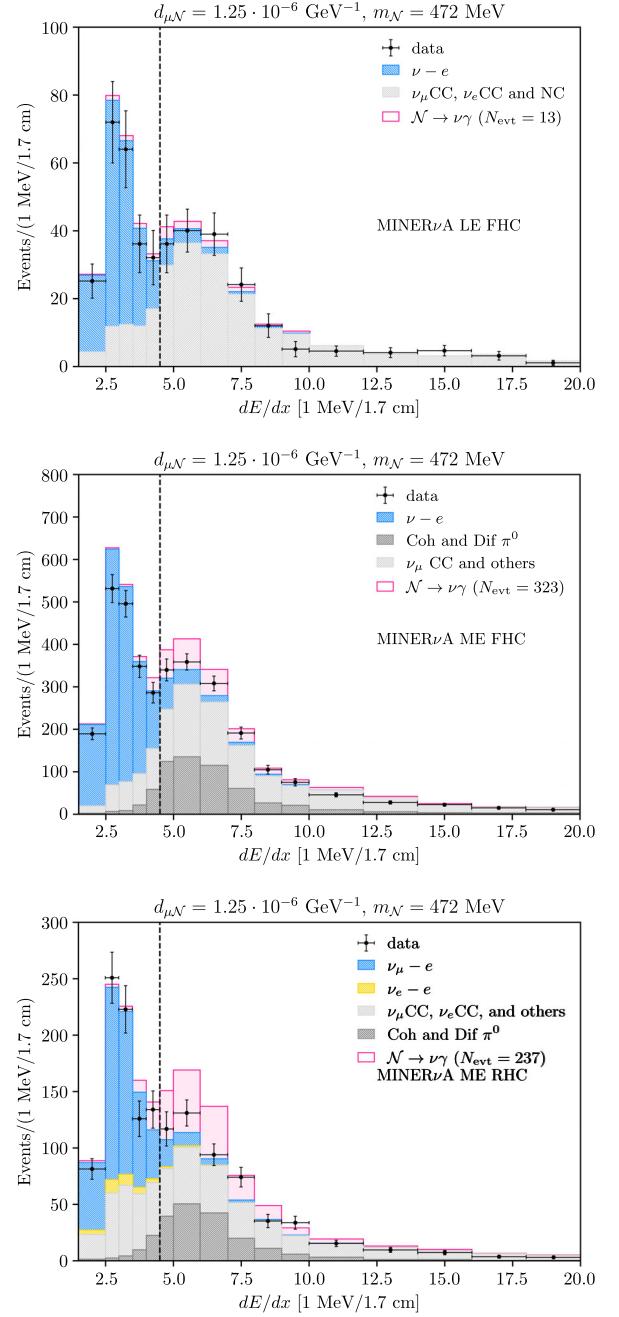


FIG. 9. The dE/dX distribution of selected events for the three MINERvA analyses. From top to bottom: LE FHC, ME FHC, and ME RHC. The latter has the largest sensitivity due to the smaller backgrounds. All distributions are shown post-MINERvA tune, except for ME RHC, where it is shown before tuning.

large acceptance due to our signal being a true single photon (as opposed to two photons from π^0 or from the e^+e^- pairs considered in the new physics model of Ref. [65]). The series of cuts are illustrated in Fig. 8, where we show the acceptance of the cuts as a function of the HNL mass. The selection acceptance is largely independent of the dipole coupling.

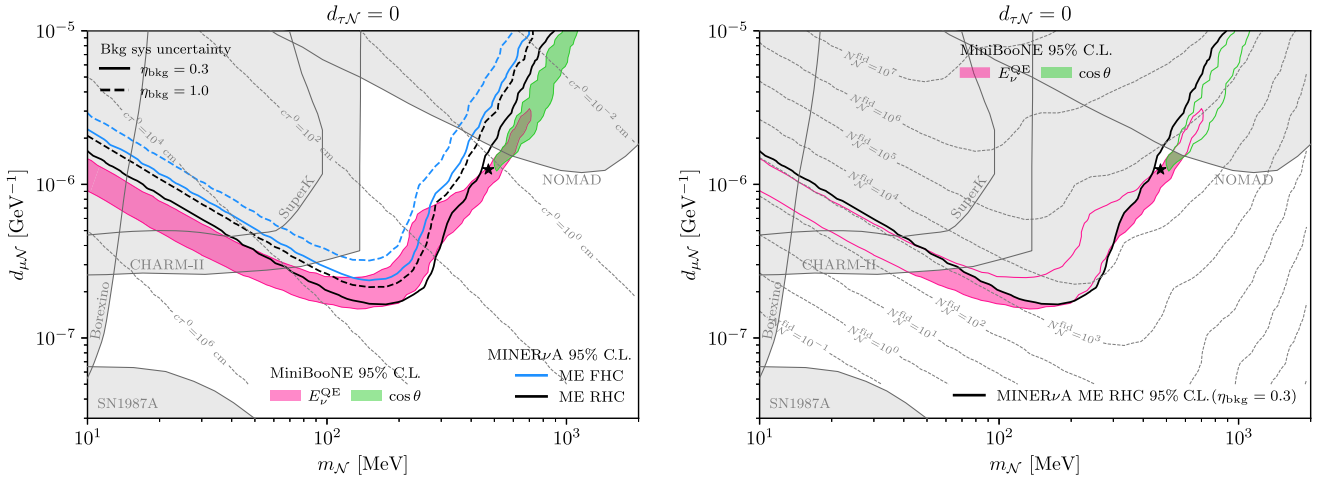


FIG. 10. (Left) MINERvA constraints in the dipole parameter space at 95% C.L. Solid lines show our nominal limits assuming a $\eta_{\text{bkg}} = 30\%$ Gaussian systematic uncertainty on the background normalization, and dashed ones show the constraints assuming an inflated uncertainty of $\eta_{\text{bkg}} = 100\%$ on the background. Regions of preference to explain MiniBooNE in the minimal dipole model are also shown as filled contours at 95% C.L. Right) Contours of constant $N_{\mathcal{N}}^{\text{fid}}$, the total number of new-physics photons that convert inside the fiducial volume, overlaid on top of the same parameter space.

We start with the cut on the reconstructed shower energy, $E_{\text{sh}} > 0.8$ GeV. To suppress ν_e CC backgrounds, a cut on the reconstructed momentum exchange under the hypothesis of neutrino-nucleon quasielastic scattering is also implemented. It is defined as $Q_{\text{QE-reco}}^2 = 2m_n(E_{\nu}^{\text{QE-reco}} - E_{\text{sh}})$, with

$$E_{\nu}^{\text{QE-reco}} = \frac{m_n E_{\text{sh}} - m_e^2/2}{m_n - E_{\text{sh}} + p_e \cos \theta_e}, \quad (6)$$

and the analysis requires $Q_{\text{QE-reco}}^2 < 0.02$ GeV². We note that for $\nu - e$ scattering this cut can be understood by the following relation,

$$Q_{\text{QE-reco}}^2 \simeq (E_{\text{sh}} \theta_{\text{sh}})^2 < \frac{m_e E_{\nu}}{2}, \quad (7)$$

where we dropped higher-order terms in electron mass and θ_{sh} . For neutrino-nucleus upscattering in the forward direction (small $\theta_{\mathcal{N}}$), two-body kinematics for an infinitely heavy nucleus gives $Q_{\text{QE-reco}}^2 \simeq (E_{\mathcal{N}} \theta_{\mathcal{N}})^2 < (E_{\nu} m_A - m_{\mathcal{N}}^2)/2$, which is a much looser constraint. In addition, the decay of \mathcal{N} introduces even more spread in the angular distribution, so we can already expect the cut on $Q_{\text{QE-reco}}^2$ to be very important.

The most stringent cut in the analysis; however, is in $E_{\text{sh}} \theta_{\text{sh}}^2$, required it to be < 3.2 MeV rad². The acceptance of this selection varies from under 10^{-4} at high masses to approximately 30% at the lowest masses. It is larger for Majorana than Dirac HNLs, as Dirac HNLs are more likely to produce backward-going photons. Finally, the last cut we

implement is the cut on the mean dE/dx of first four scintillator planes, requiring $dE/dx > 4.5$ MeV/1.7 cm. Figure 9 shows the dE/dx distributions of the SM background and HNL decay events, as well as a dashed line representing the dE/dx cut. The acceptance of this cut is the least-understood aspect of our analysis as we do not have access to a full detector simulation. The shape of our $dE/dx > 4.5$ MeV/1.7 cm distribution is assumed to be identical to that of the coherent and diffractive π^0 backgrounds shown in Fig. 9. In this approximation, we find this cut has an acceptance of 93% and 91% for FHC and RHC modes, respectively.

C. Results

The final 95% C.L. constraints on the dipole model are shown in Fig. 10. The MiniBooNE regions of preference are also shown for comparison. At the lowest values of $m_{\mathcal{N}}$ where the HNLs are long lived, our constraints are less sensitive to the MiniBooNE best-fit region than at higher masses due to the larger HNL boost factors at MINERvA as well as the smaller fiducial volume when compared to MiniBooNE. For lifetimes longer than $c\tau^0 \sim 100$ cm, the event rate in both experiments is dominated by dirt upscattering. For shorter lifetimes, the event rate is dominated by upscattering within the detector itself.

One can see that the constraints from MINERvA begin to rule out disfavor regions of parameter space preferred by MiniBooNE. However, the strongest MINERvA 2σ C.L. limits presented, which come from the ME RHC measurement and assume 30% uncertainty on the background normalization, do not rule out the intersection of the 2σ C.L. preferred regions from the MiniBooNE E_{ν}^{QE} and $\cos \theta$

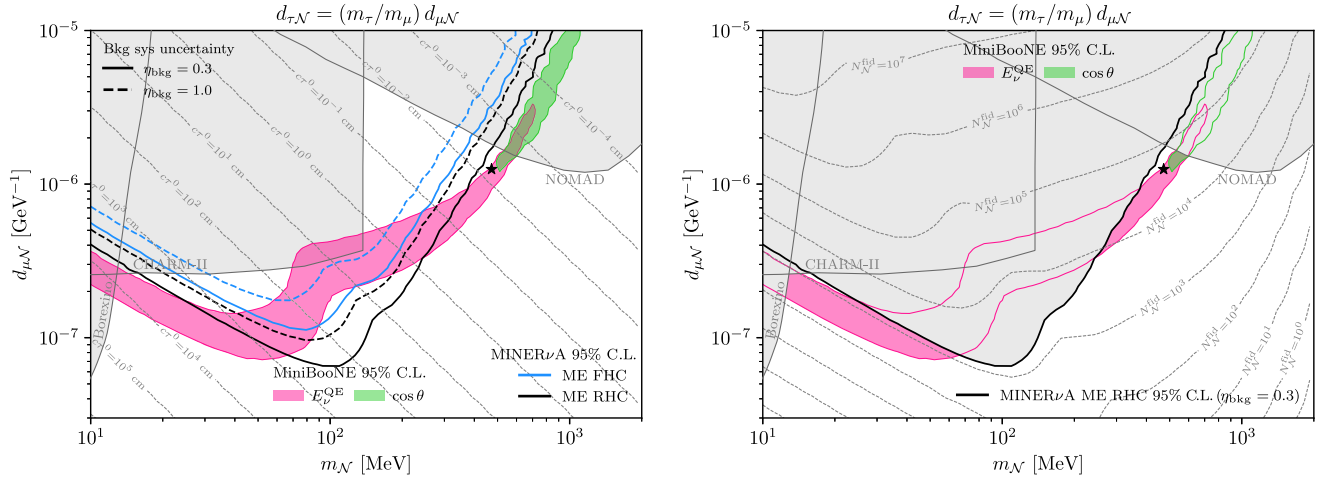


FIG. 11. Same as Fig. 10 but for the case where HNLs have a larger tau-neutrino dipole following an approximate scaling of UV completions of the operator in Eq. (1), $d_{\tau\mathcal{N}} = m_\tau/m_\mu \times d_{\mu\mathcal{N}}$. The HNL is shorter-lived due to the additional decay $\mathcal{N} \rightarrow \nu_\tau\gamma$.

distributions. This is because the $\mathcal{N} \rightarrow \nu\gamma$ acceptance in the MINERvA ES analysis decreases rapidly for larger HNL masses, as shown in Fig. 8. In the right panel of Fig. 10, we show contours of constant event rate from dipole-coupled HNL decays in MINERvA. A dedicated single-shower analysis improving the acceptance for larger HNL masses would likely be sensitive to the entire region of parameter space preferred by MiniBooNE.

In the left panel of Fig. 10, we also show conservative constraints on this model assuming 100% uncertainty on the background normalization. This is meant to address the large scale factors (up to factors of ~ 2) which have been applied to the high dE/dx backgrounds in the official MINERvA analysis [27]. These scale factors come from a tuning procedure in kinematic sideband regions, a process that could potentially wash out any signal from neutrissimo decays. An optimal analysis would perform a joint fit to both neutrissimo decays and SM high dE/dx backgrounds to derive constraints (and potentially allowed regions) on the neutrissimo model presented here; however, such an analysis is out of the scope of this paper. We also note that the MINERvA analysis does not include single-photon backgrounds such as radiative $\Delta(1232)$ decays and coherent single photons. These components are expected to be small in the energy region of $E_{\text{sh}} > 800$ MeV [66,67], but their inclusion can only make our limits stronger.

We also point out Ref. [68], in which the MINERvA Collaboration investigated an excess in the high dE/dx sideband region of a ν_e charged-current quasielastic scattering sample. Using topological variables related to the shower structure, MINERvA concluded that the excess looked more like coherent or diffractive π^0 production than single photons. This might suggest that the scale factors in Ref. [27] could also be attributed to additional π^0 events. However, the analysis presented here suggests that MINERvA may have unique sensitivity to a

neutrissimo-based explanation of the MiniBooNE excess, thus motivating a more careful separation of one and two photon events in the high dE/dx region of the MINERvA elastic scattering samples.

We also derive constraints considering nonzero $d_{\tau\mathcal{N}}$ in Figs. 11 and 12. As expected, this impacts the constraints most significantly at lower HNL masses. For both $d_{\tau\mathcal{N}} = (m_\tau/m_\mu)d_{\mu\mathcal{N}}$ and $d_{\tau\mathcal{N}} = (100 \text{ TeV})^{-1}$ The MINERvA constraints rule out a large chunk in the middle of the region preferred by the MiniBooNE E_ν^{QE} distribution. The constraints do not change for $m_{\mathcal{N}} \gtrsim 200$ MeV, thus the 2σ overlap between between the MiniBooNE E_ν^{QE} and $\cos\theta$ distributions remains valid.

VI. DISCUSSION

A number of other existing and planned neutrino experiments are sensitive to an MeV-scale dipole-coupled HNL. Super-Kamiokande can look for single photon decays from atmospheric neutrinos which upscatter into HNLs within the Earth [69]. Similarly, one can look for single-photon decays from neutrinos that upscatter into HNLs within the Earth and propagate to a large-scale terrestrial detector such as Super-Kamiokande or Borexino [70]. Constraints from these searches are sensitive to longer-lived HNLs, with typical masses $\lesssim 10$ MeV ($\lesssim 100$ MeV) in the solar (atmospheric) case. These constraints no longer apply for the two cases of nonzero $d_{\tau\mathcal{N}}$ which we consider in this work, as the HNL lifetime will be too short to reach the detector. Observed neutrino interactions from Supernova 1987A can also be used to derive constraints on the dipole model, as significant upscattering would enhance the stellar cooling rate, decreasing the neutrino flux observed on Earth [18]. These constraints require the HNL to be sufficiently long-lived that it can escape the stellar environment; thus, bounds

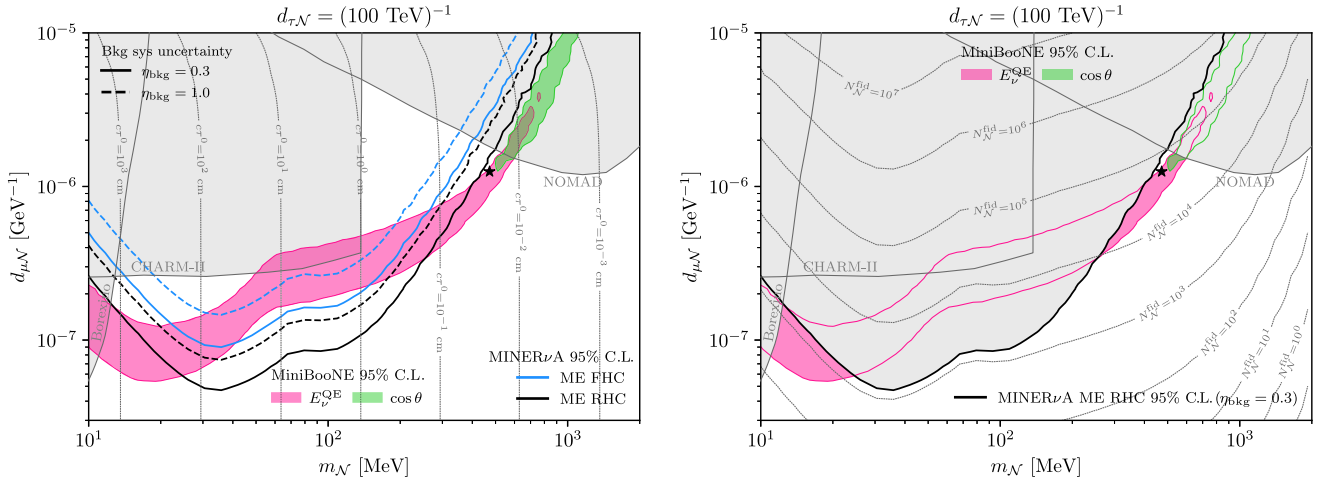


FIG. 12. Same as Fig. 10 but for the case where HNLs have a larger and fixed tau-neutrino dipole, $d_{\tau\mathcal{N}} = (100 \text{ TeV})^{-1}$. The HNL is shorter lived due to the additional decay $\mathcal{N} \rightarrow \nu_\tau \gamma$.

from Supernova 1987A also do not apply when we consider nonzero $d_{\tau\mathcal{N}}$.

We also show constraints derived from the NOMAD search for neutrino-induced single photons, recast as bounds in this parameter space in Ref. [53]. We rescale them, however, by an overall factor of $\sqrt{2}$ so as to reflect the decay rate we derived in Eq. (5). In addition, we note that these limits have been obtained with a much less sophisticated simulation than the ones performed here and that the signal was derived using only events in the preshower detector of NOMAD. In Ref. [53], it is suggested that stronger limits could be obtained by considering upscattering locations beyond the preshower detector and extended detector volumes where the HNL decay could take place.

We also include the limit imposed by CHARM-II as derived in Ref. [71]; however, we note that it was obtained with a simplified procedure. The experimental precision on the total neutrino-electron upscattering ($\nu e \rightarrow \mathcal{N}e$) cross sections. This is a reasonable assumption at low values of $m_{\mathcal{N}}$, but potentially breaks down at values close to the threshold due to differences in kinematics. A similar constraint can be set using the LSND elastic scattering measurement [18], though it is, in general, less sensitive than the CHARM-II measurement. A robust reevaluation of dipole model constraints from these electron scattering measurements is out of the scope of this paper. In addition to scattering on electrons, CHARM-II can provide new limits in the region of interest by considering coherent neutrino-nucleus upscattering followed by HNL decays into single photons [65]. The sideband with large dE/dX and large values of $E\theta^2$ can be used to set limits, as proposed in [65], although we do not expect them to be as sensitive due to larger backgrounds and larger boosts.

One can also derive constraints in the dipole-coupled HNLs from LEP through the $e^+e^- \rightarrow \mathcal{N}\nu_\ell$, which can proceed through either the γ or Z mediators [72]. However, these constraints require a strong enhancement of the mixing between the HNL and SM neutrino; as we consider such a mixing to be negligible in this model, we do not include constraints from LEP in our results.

We now discuss the potential for future constraints on dipole-coupled HNLs from planned measurements. Just like the MINERvA constraint derived in this work, a neutrino elastic scattering measurement from the NOvA experiment would be sensitive to the dipole model [73]. This is especially true of the DUNE experiment, which has the potential to make a high-statistics neutrino-electron scattering measurement [74]. This would be particularly advantageous for the THEIA@DUNE configuration [75] due to its low threshold and large volume. Dedicated searches at neutrino experiments can further improve sensitivity to this model. As discussed above, a MINERvA single-shower analysis without a stringent $E\theta^2$ cut could set much stronger constraints. Experiments which measure CEvNS, such as COHERENT, NUCLEUS, and Coherent CAPTAIN-Mills, would also be sensitive to the dipole model by looking for the coincidence of nuclear recoil from Primakoff upscattering and a single photon from the HNL decay [76]. These experiments would be most sensitive to lower mass HNLs with $m_{\mathcal{N}} \lesssim 10$ MeV due to the lower energy of typical neutrino sources for CEvNS experiments. Existing and upcoming short baseline neutrino experiments, including MicroBooNE and SBND, have the potential to be sensitive to the neutrissimo model presented in this paper through a dedicated search for single photon events [18]. Additionally, neutrino telescopes like IceCube and KM3NeT could perform searches for events with a double-bang topology from the upscattering and decay of the HNL, reaching sensitivities of $d_{\mu\mathcal{N}} \sim 10^{-7} \text{ GeV}^{-1}$ for $m_{\mathcal{N}} \lesssim 1 \text{ GeV}$ [71].

High-energy astrophysical tau neutrino observatories such as TAMBO might also be sensitive to HNL decays from ν_τ upscattering [77]. Projections for DUNE [78,79] estimate that, in the absence of backgrounds, a search for events with a double-bang morphology could reach $d_{\mu\mathcal{N}}$ values as low as $\mathcal{O}(10^{-8} \text{ GeV}^{-1})$. Finally, nuclear emulsion and liquid argon detectors at a future LHC Forward Physics Facility will also be sensitive to transition magnetic moments between HNLs and SM neutrinos [80].

VII. CONCLUSION

In this work, we have explored a mixed model comprising an eV-scale sterile neutrino and an MeV-scale dipole-coupled HNL. The former facilitates oscillations at short baselines, while the latter introduces the interactions shown in Fig. 1. The dipole-coupled HNL provides an alternative explanation of the MiniBooNE excess to the eV-scale sterile neutrino. Thus one can remove MiniBooNE from global $3 + 1$ fits, reducing tension between appearance and disappearance experiments while retaining an explanation of the LSND anomaly [23]. We take the result of the MiniBooNE-less $3 + 1$ global fit as the oscillation contribution to the MiniBooNE excess and attribute the remaining excess to decays of the dipole-coupled HNL. We find that spectral fits to the E_{ν}^{QE} and $\cos\theta$ distributions prefer different regions of parameter space in general, though solutions exist which are compatible with both distributions at the 2σ confidence level.

We have also derived constraints on the dipole-coupled HNL model using a MINERvA neutrino-electron elastic scattering measurements [26,27,61]. We find that the most sensitive $\nu - e$ scattering constraints are those obtained with the NuMI medium-energy mode in antineutrino-enhanced beam configuration. The constraints from antineutrino-mode are especially strong due to a reduction in backgrounds at high dE/dx , where we expect HNL decays to contribute. As shown in Fig. 10, MINERvA can exclude large regions of parameter space preferred by MiniBooNE, but it does not fully exclude it. There are still allowed MiniBooNE regions at the 2σ confidence level. The first is at small $m_{\mathcal{N}}$ values, where HNLs are long-lived and MINERvA's small fiducial volume and larger energies reduce the sensitivity. The second is at larger HNL masses, where the stringent $E\theta^2$ cuts reject most new physics events where decay photons tend to have larger $E\theta^2$. We note that a dedicated search at MINERvA using the same fiducial volume could significantly improve the signal efficiency in this large-mass region, and would likely have much better sensitivity, and potentially probe the entire MiniBooNE-preferred region, as shown in the right panel of Fig. 10. Nevertheless, as it stands, this mixed model of oscillations and decay is not dead yet.

ACKNOWLEDGMENTS

M. H. S. is supported by NSF Grant No. PHY-1707971. NSF Grant No. PHY-1801996 supported C. A. A., J. M. C., A. D., and N. W. K. for this work. Additionally, C. A. A. is supported by the Faculty of Arts and Sciences of Harvard University and the Alfred P. Sloan Foundation. N. W. K. is supported by the NSF Graduate Research Fellowship under Grant No. 1745302. M. A. U. is supported by the Department of Physics at the University of Cambridge and S. V. is supported by the STFC. M. H. was supported by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research, Innovation, and Science.

APPENDIX: UPSCATTERING CROSS SECTION

The cross section for \mathcal{N} production by neutrino upscattering, $\nu_\alpha + A \rightarrow \mathcal{N} + A$, has been computed several times in the literature [41,81,82]. We note, however, that all expressions we could find do not take into account the polarization of the outgoing HNL. This effect is only important when $m_{\mathcal{N}}/E$ becomes appreciably large and we find it to be a marginal effect in our calculations.

In the massless limit, a beam of left-handed polarized neutrinos will always upscatter to right-handed polarized HNLs, assuming the process takes place purely via the transition magnetic moment. This follows from the chiral structure of the vertex, $\bar{\nu}_L \sigma^{\mu\nu} \mathcal{N}_R$. However, in the massive case, spin and helicity are not equivalent, and both helicity states of \mathcal{N} can be produced. The helicity-flipping channel, $\nu_h \rightarrow \mathcal{N}_{-h}$, typically dominates, while the helicity-conserving case, $\nu_h \rightarrow \mathcal{N}_h$, will be suppressed by powers of $m_{\mathcal{N}}/E$, vanishing in the massless limit. Here $h = \pm 1$ denotes the particle's helicity and E is the typical energy scale of the scattering process.

We have calculated both terms using the `DarkNews` code [83], and show our results in Fig. 14. We show a comparison of the upscattering cross section on Carbon-12 for a few choices of HNL masses for both coherent and proton-elastic scattering regimes. Scattering on neutrons proceeds only via the neutron magnetic moment and is much smaller. We only include the proton-elastic contribution for MiniBooNE, where the proton would be invisible. This is a conservative approach when deriving the MINERvA limits.

We also show the ratio between helicity-flipping and helicity-conserving upscattering events inside the MiniBooNE detector as a function of $m_{\mathcal{N}}$ in Fig. 15. The helicity-conserving part is a small correction, except at the very largest HNL masses, where the rate is significantly smaller due to the large energy threshold for upscattering.

Several model-independent nuclear form factor parametrizations can be found in the literature. One of them is




FIG. 13. Comparison of the different nuclear form factors for ^{12}C commonly used in the literature. The dipole form factor (dashed blue) significantly overestimates the form factor at large values of the momentum exchange Q . In this paper, we use the Fourier-Bessel parametrization (solid green) for nuclei for which nuclear data is available, otherwise, we implement the Fermi-symmetrized (FS) Woods-Saxon parametrization (dotted yellow).




FIG. 14. Comparison between helicity-flipping and helicity-conserving cross sections for coherent neutrino upscattering on Carbon (left) and free protons (right) for multiple values of the HNL mass $m_{\mathcal{N}}$.

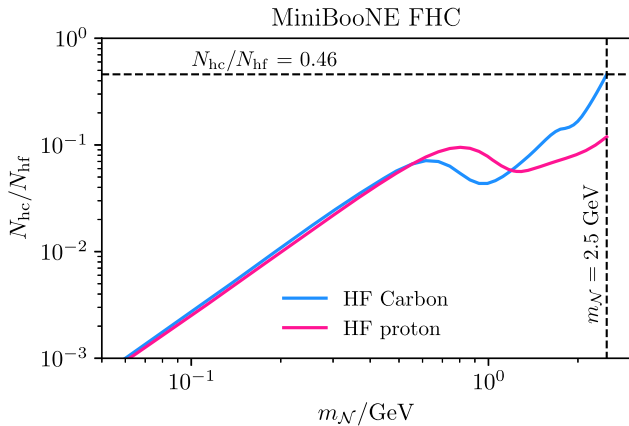


FIG. 15. The ratio between helicity-conserving and helicity-flipping upscattering events on Carbon at MiniBooNE.

the Fourier-Bessel parametrization, which models the charge density in the nucleus as a series of Bessel functions with a radial cutoff of R .

$$F^{\text{FB}}(Q) = N \times \frac{\sin(QR)}{QR} \sum_n \frac{(-1)^n a_n}{n^2 \pi^2 - Q^2}, \quad (\text{A1})$$

where $1/N = \sum_n (-1)^n a_n / n^2 \pi^2$ is a normalization factor, ensuring $F(0) = 1$. The coefficients a_n can be obtained from experimental data, which is available for a series of common nuclei [48–50]. We make use of the machine-readable files provided by Ref. [51].

For nuclei where the nuclear data cannot be found, we implement a Fermi-symmetrized Woods-Saxon form factor,

$$F^{\text{FS-WS}}(Q) = \frac{3\pi a}{r_0^2 + \pi^2 a^2} \frac{a\pi \cotanh(\pi Q a) \sin(Q r_0) - r_0 \cos(Q r_0)}{Q r_0 \sinh(\pi Q a)}, \quad (\text{A2})$$

where $a = 0.523$ fm, $r_0 = 1.03 \times A^{1/3}$ fm. These form factors correctly describe the finite nuclear radius and lead to a strong suppression of coherent scattering for $Q \gtrsim 200$ GeV. They should be contrasted with the simpler dipole parametrization

$$F^{\text{Dip}}(Q) = \frac{1}{1 + \frac{Q^2}{M_D^2}}. \quad (\text{A3})$$

used in Ref. [23], with $M_D = 1.18 + 0.83 * A^{1/3}$, and the Helmz form factor

$$F^{\text{Helmz}}(Q) = \frac{3|j_1(QR)|}{QR} e^{-Q^2 s^2/2}, \quad (\text{A4})$$

with $a = 0.523$ fm, $s = 0.9$ fm, and $R = 3.9$ fm.

We provide a comparison of the aforementioned nuclear form factors in Fig. 13. It is evident that the dipole parametrization overestimates the cross section at large values of momentum exchange Q . The more sophisticated form factors used in this work produce more forward angular distributions at MiniBooNE than what was found in the previous study of Ref. [23].

-
- [1] Q. R. Ahmad *et al.* (SNO Collaboration), *Phys. Rev. Lett.* **89**, 011302 (2002).
- [2] Y. Ashie *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. D* **71**, 112005 (2005).
- [3] 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) 071.
- [4] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou, *J. High Energy Phys.* **09** (2020) 178.
- [5] F. Capozzi, E. Di Valentino, E. Lisi, A. Marrone, A. Melchiorri, and A. Palazzo, *Phys. Rev. D* **104**, 083031 (2021).
- [6] A. A. Aguilar-Arevalo *et al.* (LSND Collaboration), *Phys. Rev. D* **64**, 112007 (2001).
- [7] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Phys. Rev. D* **103**, 052002 (2021).
- [8] M. A. Acero *et al.*, Report No. FERMILAB-PUB-22-318-ND-SCD-T, 2022.
- [9] W. Loinaz, N. Okamura, S. Rayyan, T. Takeuchi, and L. C. R. Wijewardhana, *Phys. Rev. D* **70**, 113004 (2004).
- [10] T. Adams *et al.* (NuSOnG Collaboration), *Int. J. Mod. Phys. A* **24**, 671 (2009).
- [11] S. N. Gninenko, *Phys. Rev. Lett.* **103**, 241802 (2009).
- [12] S. N. Gninenko, *Phys. Rev. D* **83**, 015015 (2011).
- [13] C. Dib, J. C. Helo, S. Kovalenko, and I. Schmidt, *Phys. Rev. D* **84**, 071301 (2011).
- [14] S. N. Gninenko, *Phys. Lett. B* **710**, 86 (2012).
- [15] M. Masip, P. Masjuan, and D. Meloni, *J. High Energy Phys.* **01** (2013) 106.
- [16] A. Radionov, *Phys. Rev. D* **88**, 015016 (2013).
- [17] P. Ballett, S. Pascoli, and M. Ross-Lonergan, *J. High Energy Phys.* **04** (2017) 102.
- [18] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, *Phys. Rev. D* **98**, 115015 (2018).
- [19] A. B. Balantekin, A. de Gouvêa, and B. Kayser, *Phys. Lett. B* **789**, 488 (2019).
- [20] S. Balaji, M. Ramirez-Quezada, and Y.-L. Zhou, *J. High Energy Phys.* **04** (2020) 178.
- [21] S. Balaji, M. Ramirez-Quezada, and Y.-L. Zhou, *J. High Energy Phys.* **12** (2020) 090.
- [22] O. Fischer, A. Hernández-Cabezudo, and T. Schwetz, *Phys. Rev. D* **101**, 075045 (2020).
- [23] S. Vergani, N. W. Kamp, A. Diaz, C. A. Argüelles, J. M. Conrad, M. H. Shaevitz, and M. A. Uchida, *Phys. Rev. D* **104**, 095005 (2021).
- [24] L. Alvarez-Ruso and E. Saul-Sala, *Eur. Phys. J. Special Topics* **230**, 4373 (2021).
- [25] J. Park *et al.* (MINERvA Collaboration), *Phys. Rev. D* **93**, 112007 (2016).
- [26] E. Valencia *et al.* (MINERvA Collaboration), *Phys. Rev. D* **100**, 092001 (2019).
- [27] L. Zazueta *et al.* (MINERvA Collaboration), *Phys. Rev. D* **107**, 012001 (2023).
- [28] R. Abbasi *et al.* (IceCube Collaboration), *Comput. Phys. Commun.* **266**, 108018 (2021).
- [29] K. Fujikawa and R. Shrock, *Phys. Rev. Lett.* **45**, 963 (1980).
- [30] P. B. Pal and L. Wolfenstein, *Phys. Rev. D* **25**, 766 (1982).
- [31] R. E. Shrock, *Nucl. Phys.* **B206**, 359 (1982).
- [32] C. A. Argüelles, N. Foppiani, and M. Hostert, *Phys. Rev. D* **105**, 095006 (2022).
- [33] M. B. Voloshin, *Sov. J. Nucl. Phys.* **48**, 512 (1988).
- [34] R. Barbieri and R. N. Mohapatra, *Phys. Lett. B* **218**, 225 (1989).
- [35] K. S. Babu and R. N. Mohapatra, *Phys. Rev. Lett.* **64**, 1705 (1990).
- [36] K. S. Babu and R. N. Mohapatra, *Phys. Rev. Lett.* **63**, 228 (1989).
- [37] M. Leurer and N. Marcus, *Phys. Lett. B* **237**, 81 (1990).
- [38] M. Lindner, B. Radovčić, and J. Welter, *J. High Energy Phys.* **07** (2017) 139.
- [39] K. S. Babu, S. Jana, and M. Lindner, *J. High Energy Phys.* **10** (2020) 040.

- [40] A. A. Aguilar-Arevalo *et al.* (MiniBooNE), *Phys. Rev. Lett.* **129**, 201801 (2022).
- [41] V. Brdar, A. Greljo, J. Kopp, and T. Opferkuch, *J. Cosmol. Astropart. Phys.* **01** (2021) 039.
- [42] V. Brdar, O. Fischer, and A. Y. Smirnov, *Phys. Rev. D* **103**, 075008 (2021).
- [43] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), [arXiv:1207.4809](https://arxiv.org/abs/1207.4809).
- [44] MiniBooNE Collaboration (private communication).
- [45] M. H. Shaevitz (MiniBooNE Collaboration), *J. Phys. Conf. Ser.* **120**, 052003 (2008).
- [46] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **599**, 28 (2009).
- [47] L. Aliaga *et al.* (MINERvA Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **743**, 130 (2014).
- [48] G. Fricke, C. Bernhardt, K. Heilig, L. A. Schaller, L. Schellenberg, E. B. Shera, and C. W. de Jager, *At. Data Nucl. Data Tables* **60**, 177 (1995).
- [49] H. De Vries, C. W. De Jager, and C. De Vries, *At. Data Nucl. Data Tables* **36**, 495 (1987).
- [50] C. W. De Jager, H. De Vries, and C. De Vries, *At. Data Nucl. Data Tables* **14**, 479 (1974); **16**, 580(E) (1975).
- [51] <http://discovery.phys.virginia.edu/research/groups/ncd/index.html>.
- [52] S. S. Wilks, *Ann. Math. Stat.* **9**, 60 (1938).
- [53] S. N. Gninenko and N. V. Krasnikov, *Phys. Lett. B* **450**, 165 (1999).
- [54] J. Dorenbosch *et al.* (CHARM Collaboration), *Z. Phys. C* **41**, 567 (1989); **51**, 142 (1991).
- [55] P. Vilain *et al.* (CHARM-II Collaboration), *Phys. Lett. B* **335**, 246 (1994).
- [56] L. B. Auerbach *et al.* (LSND Collaboration), *Phys. Rev. D* **63**, 112001 (2001).
- [57] M. Deniz *et al.* (TEXONO Collaboration), *Phys. Rev. D* **81**, 072001 (2010).
- [58] G. Bellini *et al.*, *Phys. Rev. Lett.* **107**, 141302 (2011).
- [59] M. Agostini *et al.* (Borexino Collaboration), *Phys. Rev. D* **100**, 082004 (2019).
- [60] S. Fukuda *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. Lett.* **86**, 5651 (2001).
- [61] J. Park, Ph.D. thesis, University of Rochester, 2013.
- [62] L. Aliaga *et al.* (MINERvA Collaboration), *Phys. Rev. D* **94**, 092005 (2016); **95**, 039903(A) (2017).
- [63] A. Bashyal, Ph.D. thesis, Oregon State University, 2021.
- [64] E. Valencia-Rodriguez, Ph.D. thesis, Guanajuato University, 2016.
- [65] C. A. Argüelles, M. Hostert, and Y.-D. Tsai, *Phys. Rev. Lett.* **123**, 261801 (2019).
- [66] D. Rein and L. M. Sehgal, *Phys. Lett.* **104B**, 394 (1981); **106B**, 513(E) (1981).
- [67] E. Wang, L. Alvarez-Ruso, and J. Nieves, *Phys. Lett. B* **740**, 16 (2015).
- [68] J. Wolcott *et al.* (MINERvA Collaboration), *Phys. Rev. Lett.* **117**, 111801 (2016).
- [69] R. A. Gustafson, R. Plestid, and I. M. Shoemaker, *Phys. Rev. D* **106**, 095037 (2022).
- [70] R. Plestid, *Phys. Rev. D* **104**, 075027 (2021).
- [71] P. Coloma, P. A. N. Machado, I. Martinez-Soler, and I. M. Shoemaker, *Phys. Rev. Lett.* **119**, 201804 (2017).
- [72] F. Delgado, L. Duarte, J. Jones-Perez, C. Manrique-Chavil, and S. Peña, *J. High Energy Phys.* **09** (2022) 079.
- [73] J. Bian, in *Proceeding for the APS DPF 2017 conference in Batavia, IL* (2017), <https://inspirehep.net/conferences/1430017?ui-citation-summary=true>.
- [74] C. M. Marshall, K. S. McFarland, and C. Wilkinson, *Phys. Rev. D* **101**, 032002 (2020).
- [75] M. Askins *et al.* (Theia Collaboration), *Eur. Phys. J. C* **80**, 416 (2020).
- [76] P. D. Bolton, F. F. Deppisch, K. Fridell, J. Harz, C. Hati, and S. Kulkarni, *Phys. Rev. D* **106**, 035036 (2022).
- [77] A. Romero-Wolf *et al.*, in *Latin American Strategy Forum for Research Infrastructure* (2020), <https://inspirehep.net/conferences/1726700?ui-citation-summary=true>.
- [78] M. Atkinson, P. Coloma, I. Martinez-Soler, N. Rocco, and I. M. Shoemaker, *J. High Energy Phys.* **04** (2022) 174.
- [79] T. Schwetz, A. Zhou, and J.-Y. Zhu, *J. High Energy Phys.* **07** (2021) 200.
- [80] A. Ismail, S. Jana, and R. M. Abraham, *Phys. Rev. D* **105**, 055008 (2022).
- [81] P. Vogel and J. Engel, *Phys. Rev. D* **39**, 3378 (1989).
- [82] R. Harnik, J. Kopp, and P. A. N. Machado, *J. Cosmol. Astropart. Phys.* **07** (2012) 026.
- [83] A. Abdullahi, J. Hoefken, M. Hostert, D. Massaro, and S. Pascoli, DarkNews, <https://github.com/mhostert/DarkNews> (2019).