Efficiently exploring multidimensional parameter spaces beyond the Standard Model

Carlos A. Argüelles,[1,*] Nicolò Foppiani ,[1,†] and Matheus Hostert [2,3,4,‡]

[1]*Department of Physics and Laboratory for Particle Physics and Cosmology,*
Harvard University, Cambridge, Massachusetts 02138, USA
[2]*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2J 2W9, Canada*
[3]*School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA*
[4]*William I. Fine Theoretical Physics Institute, School of Physics and Astronomy,*
University of Minnesota, Minneapolis, Minnesota 55455, USA

We propose a method to ease the challenges of exploring multidimensional parameter spaces in beyond-the-Standard Model theories. We evaluate the model likelihood for any choice of parameters by sampling the theory parameters intelligently and building a kernel density estimator. By reducing the number of expensive Monte Carlo simulations, this method provides an efficient way to test complex theories. We illustrate our technique to set new limits on a short-lived heavy neutrino N, proposed as an explanation of anomalies in neutrino experiments. Using a search for lepton pairs in the T2K near detector, we find exclusion limits on the model parameters in a vast region of parameter space, fully exploiting the advantages of our new method. With a single Monte Carlo simulation, we obtain the differential event rate for arbitrary choices of model parameters, allowing us to cast limits on any slice of the model parameter space. We conclude that N particles with lifetimes greater than $c\tau^0 \gtrsim 3$ cm are excluded by T2K data. We also derive model-independent constraints in terms of the total rate, lifetime, and N mass and provide an approximated analytical formula. This method can be applied in other branches of physics to explore the landscape of theory parameters efficiently.

I. INTRODUCTION

Neutrinos are a unique ingredient of the Standard Model (SM). The absence of electric charge and their extremely small but nonvanishing mass implies that, contrary to all other fermions, neutrinos do not have their properties uniquely determined by the SM gauge group, $G = \mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$. Indeed, to uniquely determine the origin of neutrino masses, we are left to choose between the existence of new symmetries, such as $U(1)_{B-L}$, or new scales, such as the Majorana mass of their right-handed partners. The latter is the route chosen in the canonical Type-I seesaw mechanism [1–9], which has long been the leading motivation for experimental searches for feebly-interacting Majorana particles in cosmological or laboratory settings [2]. As a

complete singlet under G, right-handed neutrinos could also provide unique insight into the possible existence of other hypothetical particles, such as dark matter, the dark photon, or additional Higgs bosons.

While neutrino experiments have achieved remarkable success in proving that neutrinos change flavor over macroscopical distances, global data still does not point to a coherent picture. Short-baseline experiments, characterized by baselines and energies of $L/E \sim 1$ km/GeV, provide significant outliers in the framework of three-neutrino oscillations. This is led mostly by the 4.8σ excess of electronlike events at MiniBooNE [3–8], and the 3.8σ excess of inverse-beta-decay events at LSND [9–13]. For a long time, solutions based on a light sterile neutrino with Δm^2_{41} of $\mathcal{O}(1)$ eV^2 were the leading candidates for a beyond-the-SM explanation but have yet to successfully overcome challenges with accelerator and reactor experiments, as well as cosmological observations [14]. For a recent review on the status of this topic, see Ref. [15]. Since then, theoretical activity in low-scale dark sectors has intensified and, unsurprisingly, new dark-sector solutions to the short-baseline puzzle have been brought to light [16–40].

Among them is the possibility that light dark particles are produced in the scattering of neutrinos with matter and

[*]carguelles@fas.harvard.edu
[†]nicolofoppiani@g.harvard.edu
[‡]mhostert@perimeterinstitute.ca

misidentified as electron-neutrinos due to their electromagnetic decays. These models have been popularized by their connection to neutrino masses in low-scale seesaw models, the possibility to explain other low-energy anomalies, and, chiefly, due to their falsifiability.

Often, these dark sector solutions involve several new particles and, with them, several independent parameters. While the experimental signatures are straightforward to identify, thoroughly exploring the model parameter space can quickly become unmanageable. This curse of dimensionality is especially burdensome due to the expensive Monte Carlo simulations of the experiments. This article provides a solution to this problem using a new method based on reweighing Monte Carlo samples. Thoroughly exploring parameter spaces is essential both when placing constraints, since parameters can often be correlated, and when searching for degenerate new-physics solutions since they can be missed in coarse scanning procedures. Our method is a generic way to explore any model with a complex parameter space; thus, it applies to several branches of physics.

The novelty of our technique lies in the usage of a single Monte Carlo simulation that simultaneously samples physical quantities, like phase space variables, as well as model parameters, such as the masses of heavy neutrinos and the dark photon. With sufficient statistics, these samples can be used to construct a kernel density estimator (KDE), which computes the model prediction and corresponding likelihood for any choice of model parameters within the boundaries of the simulation. The result is a fast interpolation of the posterior probability of the model and greater flexibility in determining confidence intervals in various slices of parameter space. This method does away with the usual procedure used in the phenomenology community of repeating the full experimental simulation on an n-dimensional grid, which can be a much more inefficient way to build the posterior. By adaptively sampling the model parameter space, we maximize the sampling around the regions of interest.

We illustrate our proposal in the context of a dark sector model that can explain the MiniBooNE excess. The model will contain heavy neutrinos and a dark photon, whose properties are defined by parameters embedded in a multidimensional parameter space. We apply our method to a search for e^+e^- in the multicomponent near detector of T2K, ND280 [41]. The signatures arise from the upscattering of neutrinos inside the high-density region of the detector, followed by the decay of the heavy neutrino into e^+e^- pairs inside the gaseous argon (GAr) time-projection chambers (TPC) of ND280. By leveraging the power of our technique, we can take advantage of our detailed detector simulation throughout a much broader parameter space.

This article is divided as follows. In Sec. II we introduce the model we study in this work, providing a phenomenology-friendly parametrization. Section III provides details

on the ND280 detector and the analyses we use to set constraints. In Sec. IV we discuss our dedicated Monte Carlo simulation and introduce our reweighing scheme and KDE techniques. Finally, in Sec. V, we present the resulting limits in slices of parameter space and conclude in Sec. VI.

II. DARK NEUTRINOS

The idea that low-scale seesaw extensions of the SM can coexist with new gauge symmetries, most famously with $B - L$ [42–51], has been discussed throughout the literature also in the context of baryonic [52], leptonic [53,54], or completely hidden gauge symmetries [55–61]. These models present a complicated mass spectrum and a self-interacting dark sector that can be challenging to identify experimentally. Nevertheless, they remain viable and testable examples of low-scale neutrino mass mechanisms and deserve experimental scrutiny.

In this article, we focus on a low-scale dark sector containing heavy neutrino states and a broken $U(1)_D$ gauge symmetry [25–31]. The heavy neutrinos interact with the mediator of the dark gauge group, the dark photon, and with SM neutrinos via mixing. Through a combination of neutrino mixing and kinetic mixing between the SM photon and the dark photon, this dark sector leads to new interactions of neutrinos with charged particles. It contains heavy neutrino states that decay primarily through the new force. This model is especially interesting in the context of short-baseline anomalies as it predicts the production of heavy neutrinos inside detectors, which subsequently decay to electromagnetic final states, mimicking $\nu_\mu \to \nu_e$ appearance signatures.

We start with the definition of a simplified, low-energy Lagrangian that will be used throughout this work. We discuss mechanisms to UV-complete this Lagrangian in Sec. II A. The minimal particle content we consider contains a single mediator, corresponding to a kinetically-mixed dark photon Z', and heavy neutrino states $\nu_{i\geq4}$. We provide further details on the UV completions of the model at the end of this section. In terms of the physical fields, our Lagrangian reads

$$\mathcal{L} \supset \mathcal{L}_{\nu\text{-mass}} + \frac{m_{Z'}^2}{2} Z'^\mu Z'_\mu + Z'_\mu (e\epsilon J_{\text{EM}}^\mu + g_D J_D^\mu), \quad (1)$$

where $\mathcal{L}_{\nu\text{-mass}}$ contains all the mass terms for the neutrino fields after proper diagonalization. The dark photon interacts with the SM electromagnetic current, J_{EM}^μ, proportionally to the electric charge e and kinetic mixing ϵ. It also interacts with the neutral leptons dark current, J_D^μ, proportionally to the gauge coupling g_D. The above Lagrangian includes all interactions of interest in the limit of small ϵ and $(m_{Z'}/m_Z)^2$.

In terms of an interaction matrix V, the dark current in the mass basis is given by

$$J_D^\mu = \sum_{i,j}^{n+3} V_{ij}\overline{\nu}_i\gamma^\mu\nu_j, \tag{2}$$

where n is the number of heavy neutrino states. Here, $\nu_{1,2,3}$ are the mostly-SM-flavor light neutrinos, and $\nu_{i\geq 4}$ are the heavy neutrinos that contain small admixtures of SM flavors. We can express V_{ij} in terms of the mixing matrix U_{di} between the mass eigenstates i and dark flavor states d as $V_{ij} = \sum_d Q_d U_{di}^* U_{dj}$, where the index d runs over all dark-neutrino flavors ν_d of $U(1)_D$ charge Q_d. Assuming $Q_d = 1$ for all flavors, $|V_{ij}| \leq 1$ for all i and j due to the unitarity of the full neutrino mixing matrix. Since we are not interested in the specifics of the flavor structure of the full model, we stick with the generic notation of Eq. (2), noting that experimental constraints require $V_{ij} \ll 1$ if either i or j are in $\{1, 2, 3\}$.

We now consider the upscattering of light neutrinos into one of the n heavy neutrinos states ν_N, hereafter referred to as N for brevity, and its subsequent decay into lighter neutrinos and an e^+e^- pair. Specifically,

$$\hat{\nu}_\mu + A \to (N \to \nu e^+ e^-) + A, \tag{3}$$

where $\hat{\nu}_\mu$ stands for the coherent superposition of $\nu_{1,2,3}$ produced in the neutrino beam, and ν for all possible daughter neutrinos. Throughout this work, we consider only coherent scattering on the nucleus A.

A. Upscattering and decay

The coherent neutrino-nucleus scattering is mediated by the dark photon with amplitude

$$\mathcal{M}_{\rm ups}^{Z'} = \frac{e\epsilon g_D}{q^2 - m_{Z'}^2}\ell_\mu h^\mu, \tag{4}$$

where q^2 is the momentum exchange with the nucleus, $h^\mu = \langle A|J_{\rm EM}^\mu|A\rangle$ is the elastic electromagnetic transition amplitude for the nuclear ground state of A, and ℓ^μ is the leptonic current

$$\ell^\mu = \langle N|J_D^\mu|\hat{\nu}_\mu\rangle = \frac{\sum_{i\leq 3} U_{\mu i}^* V_{iN}\langle N|\overline{N}\gamma^\mu\nu_i|\nu_i\rangle}{(\sum_{k\leq 3}|U_{ki}|^2)^{1/2}},$$
$$= V_{\mu N}\langle N|\overline{N}\gamma^\mu\nu_i|\nu_i\rangle, \tag{5}$$

where we defined the vertex factor

$$V_{\alpha N} \equiv \frac{\sum_{i\leq 3} U_{\alpha i}^* V_{iN}}{(\sum_{k\leq 3}|U_{ki}|^2)^{1/2}}. \tag{6}$$

In a model with a single dark flavor $d = D$ and one heavy neutrino $N = \nu_4$, one can show that $V_{\alpha N} = U_{\alpha 4}|U_{D4}|^2 \simeq U_{\alpha 4}$, which is small and directly constrained by laboratory experiments. The full cross section is then computed in the usual fashion. We have implemented a data-driven Fourier-Bessel parametrization for the nuclear form factors [62]. An approximate analytic formula for the full cross section is provided in Sec. II B.

The decay process can be computed similarly, now summing over the daughter neutrinos incoherently,

$$|\mathcal{M}_{\rm dec}^{Z'}|^2 \equiv \sum_{i<N}|V_{Ni}\mathcal{M}(m_i)|^2 \simeq |\mathcal{M}(0)|^2\sum_{i<N}|V_{Ni}|^2, \tag{7}$$

where we factorize the matrix elements assuming that all daughter neutrinos have a negligible mass with respect to m_N. We define the remaining vertex factor as

$$|V_N|^2 = \sum_{i<N}|V_{iN}|^2. \tag{8}$$

As before, for a model with a single dark flavor and one heavy neutrino, $|V_N|^2 = |U_{D4}|^2(1 - |U_{D4}|^2) \simeq |U_{e4}|^2 + |U_{\mu 4}|^2 + |U_{\tau 4}|^2$. Clearly, $|V_N|^2$ may be similar in size or much larger than $|V_{\mu N}|^2$, depending on the flavor structure of the model.[1] This way, the production cross section is effectively decoupled from the lifetime of N. This statement is irrelevant for light dark photons ($m_{Z'} < m_N$), as the decay is always prompt, rendering most signatures independent of $|V_N|^2$. However, in MiniBooNE explanations where N decays via an off shell dark photon, the requirement $|V_N| > |V_{\mu N}|$ helps ensure that the production of N, as well as its decays, happen inside the detector.

We also keep the number of daughter neutrinos unspecified to effectively cover models where N does not decay only into $\nu_{1,2,3}$, but also into other heavy neutrinos ν_j with $3 < j < N$. In this case, V_{jN} is mainly insensitive to the direct limits on the mixing of active and heavy neutrinos, $|U_{\alpha 4}|^2$, and can be of order one. The properties of these new states are rather model-dependent, so we conservatively consider them invisible and not observable. For simplicity, we require that ν_j be light enough such that $(m_N - m_j)/m_N \ll 1$. Therefore, we do not consider scenarios with small mass splittings between the upscattered and the daughter neutrinos.

The relevant decay rate for a Dirac N with off shell Z' is

$$\Gamma_{N\to\nu e^+ e^-} = \frac{\alpha\alpha_D\epsilon^2|V_N|^2}{48\pi}\frac{m_N^5}{m_{Z'}^4}L(m_N^2/m_{Z'}^2), \tag{9}$$

where $L(x) = \frac{12}{x^4}(x - \frac{x^2}{2} - \frac{x^3}{6} - (1-x)\log\frac{1}{1-x})$, with $L(0) = 1$. For a light, on shell Z' we need only compute $N \to \nu Z'$ since $Z' \to e^+e^-$ is always prompt,

$$\Gamma_{N\to\nu Z'} = \frac{\alpha_D|V_N|^2}{4}\frac{m_N^3}{m_{Z'}^2}\left(1 - \frac{m_{Z'}^2}{m_N^2}\right)^2\left(\frac{1}{2} + \frac{m_{Z'}^2}{m_N^2}\right). \tag{10}$$

Note that the decay rate is bounded from above and below by

$$\Gamma_N(|V_N| = |V_{\mu N}|) < \Gamma_N < \Gamma_N(|V_N| = 1), \tag{11}$$

---

[1]This was the idea proposed in Ref. [27], where by virtue of $|U_{\tau 4}|^2 \gg |U_{\mu 4}|^2$, the daughter neutrino produced had a significant admixture of the tau flavor.

where MiniBooNE explanations prefer to saturate the right-most inequality whenever $m_N < m_{Z'}$. While for the light Z' case the lifetimes of both N and Z' are always prompt, for the heavy case they are much longer. For instance, taking $|V_N| = 1$ and $m_{Z'} = 1.25$ GeV, we find

$$c\tau_{\min}^0 \simeq 1 \text{ cm} \times \left(\frac{10^{-2}}{\epsilon}\right)^2 \left(\frac{100 \text{ MeV}}{m_N}\right)^5 \left(\frac{m_{Z'}}{1.25 \text{ GeV}}\right)^4. \quad (12)$$

A diagram of the upscattering and decay processes is shown in Fig. 1, including the dependence on the new physics couplings at each vertex.

B. Benchmark points and MiniBooNE

We now comment on the broader context of the MiniBooNE anomaly and present two choices of model parameters that will help us benchmark the MiniBooNE explanation in the context of dark photon models. Despite being around for over two decades, the MiniBooNE and LSND anomalies have remained unsolved. Most recently, the origin of the MiniBooNE excess was recently searched for in the MicroBooNE experiment [63–66]. As argued in Ref. [67], the sterile neutrino interpretation, although disfavored, is not entirely ruled out, and several hypotheses behind the excess remain untested, including those involving large ν_e disappearance [68]. This observation is also true for explanations that rely on particle misidentification, such as models with $e^\pm \leftrightarrow \gamma$ or $e^\pm \leftrightarrow e^+e^-$ misidentification. Several models exploring this mechanism have been put forward. However, only a small subset of those can also explain LSND [40] due to the much-harder-to-fake inverse-beta-decay signature. In that case, a neutral mediator in neutrino-nucleus scattering processes should kick out a neutron from inside the carbon nucleus. Not only is this a negligible effect for a dark photon mediator, but it also requires larger neutrino energies to produce the heavy neutrinos and remain above the $E_e > 20$ MeV analysis threshold. Therefore, we proceed to present benchmark points compatible with the MiniBooNE observation only. Therefore, we proceed to present benchmark points compatible with the MiniBooNE observation only (see Table I).

a. Benchmark A, light Z': Previously, in Ref. [28], we picked the same benchmark as in Ref. [25] to present bounds set by the neutrino-electron elastic scattering measurement performed by MINERνA. Here we would like to target the low-N mass region of the parameter space. This part of the parameter space is not constrained very effectively by MINERνA due to systematic uncertainties in the background estimate.

b. Benchmark B, heavy Z': This benchmark is inspired by the benchmarks provided in Refs. [29,31]. It illustrates the case of a heavy dark photon, where the coherent upscattering contribution is not dominant but still significant.

Both choices of parameters above can explain the MiniBooNE energy spectrum but not the angular spectrum. The exchange of a dark photon with the nucleus gives rise
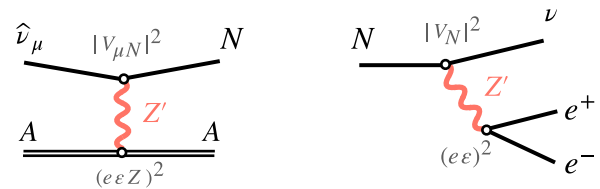

FIG. 1. The diagrams for coherent neutrino-nucleus upscattering ($\hat{\nu}_\mu A \to NA$) and heavy neutrino decays ($N \to \nu e^+e^-$) are considered in this work. We indicate the relevant parametrization of each interaction vertex.

to very low-Q^2 processes and, therefore, very forward e^+e^- final states, making the model in apparent contradiction with the MiniBooNE observation. Quantifying this tension, however, is not currently possible due to the lack of public information on the background systematics in $\cos\theta$. Since systematic uncertainties in the background prediction dominate the significance of the MiniBooNE excess, a proper fit should include the systematic uncertainties in the angular bins and, most importantly, their correlations.

Models with helicity-flipping interactions and heavy mediators, such as scalar mediator models [40], have a better chance of describing the angular spectrum. These also have the advantage of having larger cross sections with neutrons and being interesting in the context of LSND. Due to their broader angular spectrum, we expect these models to have smaller selection efficiencies at T2K than the ones we find for the dark photon model. We leave the exploration of these models for future work after proper fits to the MiniBooNE excess have been performed.

It should also be noted that there have been a series of constraints posed on the model above, from accelerator
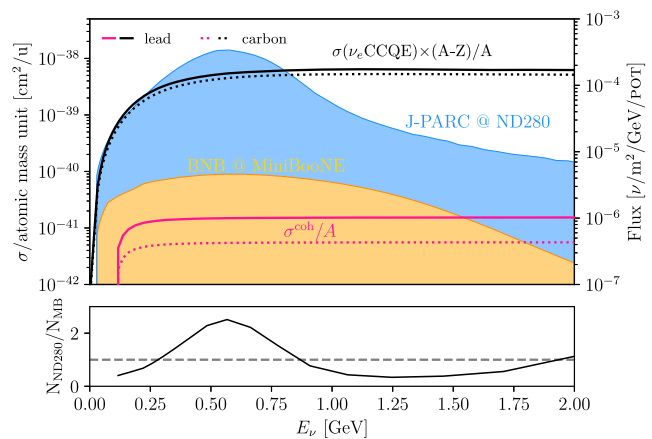

FIG. 2. Comparison between relevant neutrino fluxes and cross sections. Although MiniBooNE has a much larger mass, ND280 benefits from a significantly larger neutrino flux and higher Z materials. The solid and dotted lines show the cross section per atomic mass unit for lead and carbon, respectively. The ν_eCCQE cross sections are shown as black lines and the coherent upscattering cross sections for our heavy dark photon benchmark (B) as pink lines. When considering all the active material in ND280 and MiniBooNE we found that the ratio of upscattering between the two experiments is $\mathcal{O}(1)$ across the energy spectrum.

neutrino experiments to kaon decays. We note the study of Ref. [69], where the authors point out a large set of experimental observables that can be used to constrain MiniBooNE explanations, among which the ND280 data that we make use of. In our analysis, we properly take into account detector effects and systematics, carefully describing the geometry of the detector, which is an essential ingredient to correctly determine the bounds on the heavy Z' case, which are extremely sensitive to the lifetime of the heavy neutrino.

III. ND280 ANALYSES

ND280 is the off-axis near detector of T2K located at 280 m from the target at an angle of $\sim 2.042°$ with respect to the beam [70]. The mean neutrino energy at this location is very similar to that of the Booster Neutrino Beam, where MiniBooNE is located. The comparison of the two fluxes in Fig. 2 clearly shows that the flux seen by ND280 is significantly larger than that seen by MiniBooNE for the same exposure. The active mass, however, is much smaller—

MiniBooNE contains a total of 818 t of mineral liquid scintillator (CH_2) compared with 18 t of total active mass. Considering these two elements, we expect a similar number of upscattering events to happen in the two detectors, enabling T2K to directly test the dark neutrino interpretation of the MiniBooNE excess.

Our analysis reinterprets two public results by T2K: the search for the in-flight decays of long-lived heavy neutrinos [41] and the ν_eCCQE cross section measurement [71]. The former directly searched for appearing e^+e^- pairs inside the low-density region of the detector, while the latter measured the rate of single photons that convert into e^+e^- pairs inside one of the tracking components of the detector. We discuss the detector components below to then discuss the two analyses.

A. The ND280 detector

ND280 is a highly segmented and magnetized detector. The detector modules used in our work are shown in Fig. 3 and constitute most of the active volume of the detector.
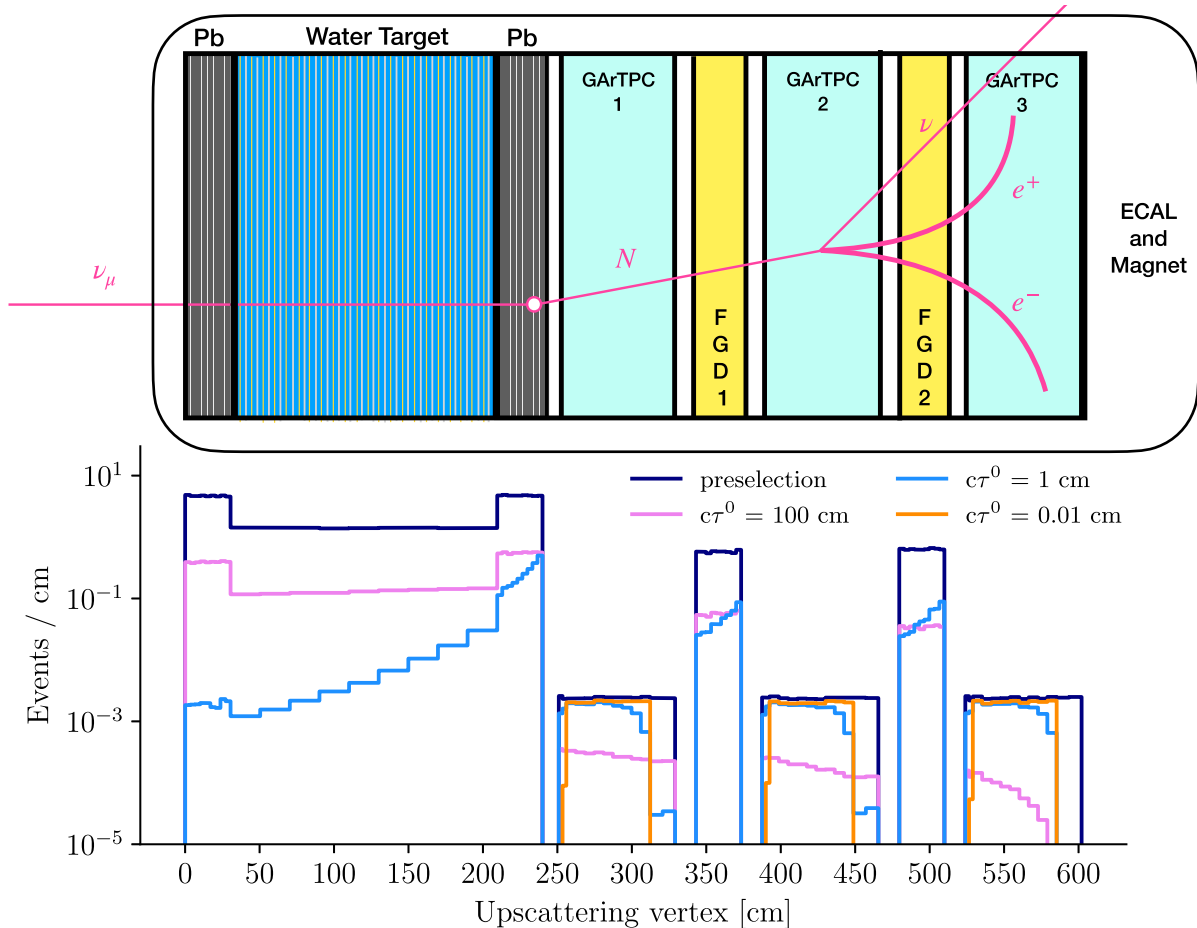

FIG. 3. Diagram of the T2K near detector, ND280, showing all the active components of the detector and the new-physics signature we are interested in. Below, we show the event rate distribution as a function of the upscattering position z before and after geometrical and analysis selection. For long lifetimes ($c\tau^0 \gtrsim 1$ cm), the event rate is dominated by upscattering on lead, while for much smaller values, only upscattering on the gaseous argon modules contributes.

TABLE I. Parameters of our two benchmark points. These choices are compatible with the excess of events observed at MiniBooNE.

Benchmark	(A) Light Z'	(B) Heavy Z'
m_N (MeV)	100	100
$m_{Z'}$ (GeV)	0.03	1.25
$|V_{\mu N}|^2$	8×10^{-9}	2.2×10^{-7}
α_D	1/4	0.4
ϵ	1.7×10^{-4}	2×10^{-2}
$|V_N|^2$	$|V_{\mu N}|^2$	1
$c\tau_N^0$ (cm)	7×10^{-5}	0.54
$N_{\text{upscattering}}^{\text{T2K}}$	15.5	560
$N_{\text{decay}}^{\text{T2K}}$	15.5	4.6

The first three modules constitute the P∅D detector, a layered arrangement of high-Z material such as lead and brass, intertwined with plastic scintillators and water bags that serves both as an electromagnetic calorimeter (ECAL) and an active water target. This subdetector is specifically designed to study π^0 production and neutrino cross sections in water, both critical inputs for the oscillation analyses using the Super-Kamiokande far detector. The first and third ECALs contain only layers of lead and scintillator plates. The module in between contains the water bags as well as layers of brass and scintillator plates. Downstream we have the tracking modules composed of three gaseous argon time-projection chambers (GArTPC), separated by fine-grained scintillator detectors (FGD). These components have fewer neutrino interactions and provide a better environment for particle identification. The whole detector is in a 0.2 T magnetic field and is surrounded by additional side and back ECAL as well as side muon detectors. A detailed description of our detector simulation is outlined in Appendix C. The dimensions and composition of each region are summarized in Table II (see Refs. [72–74] for more information on ND280). In our simulation, we assume that the lead, brass, water, and scintillator layers are distributed uniformly inside each module.

Each GArTPC is enclosed in a plastic and aluminum cage. The cage, in turn, is composed of an external wall (10.1 kg), an internal volume of CO_2 gas, and an internal wall (6.9 kg). Inside the TPC, on the Y-Z plane, one can find the cathode (6 kg, made of 34% C, 25% O, 17% Cu, 17% Si). Even though these materials are the closest to the active volume of the TPC, we neglect them as the total number of neutrino interactions recorded in them of order thousand times smaller than in the other targets.

B. Analysis I: Heavy neutrino searches in the GAr TPC

The analysis in [41] looked for the decay in flight (DIF) of heavy neutrinos inside the three GAr TPCs. Heavy neutrinos are produced in the target through mixing between active and sterile neutrinos, and, after propagating from the target to the detector, they decay in the detector TPCs. They look for multiple final states, although the relevant one for our analysis is $N \to \nu e^+ e^-$, which also extends to the lowest masses, with a threshold of $m_N \sim 1$ MeV. This analysis benefits from a clear signature with zero background, as the total argon mass is so tiny that neutrino interactions inside the TPC do not produce a relevant background. In order to achieve the zero background, the original selection imposes a tight fiducial volume cut in the TPC, with a requirement of no additional visible energy deposition in the detector in addition to the charged particles produced in the TPC. Our model can be tested with this analysis because it predicts a large coherent cross section, resulting in a very low-energy nuclear recoil that is invisible in these detectors.

This analysis is a counting experiment performed over 12.34×10^{20} proton on target (POT) in neutrino mode and 6.29×10^{20} POT in antineutrino mode. Having observed zero $e^+ e^-$ events over the neutrino background expectation of 0.563 in neutrino mode and 0.015 in antineutrino mode, the analysis sets strong limits on new physics. Limits on long-lived heavy neutrinos produced in kaon decays at the target have been discussed in Ref. [41,75].

For dark neutrino models, those can be recast by considering the production of N via upscattering inside

TABLE II. The active and fiducial volume dimensions, mass, and composition of each ND280 module simulated in our analysis. No fiducial volume is shown for the P∅D because we do not employ this detector for measuring heavy neutrino decays but only for its production. For the GAr modules, we show the fiducial volume and mass of each TPC, taking $\rho_{\text{GAr}} = 1.78$ g/cm^3.

| ND280 module	Active volume $X \times Y \times Z$ [cm]	$z_{\text{begin}}^{\text{active}}$	$M_{\text{tot}}^{\text{active}}$	Fiducial volume $X \times Y \times Z$ [cm]	$z_{\text{begin}}^{\text{fiducial}}$	$M_{\text{tot}}^{\text{fiducial}}$	composition (in mass)
P0D-ECAL1	$210 \times 224 \times 30.5$	0	2.9				6.5% H, 40% C, 53.5% Pb
P0D-water	$210 \times 224 \times 179$	30.5	10				10% H, 43% C, 22% O, 16% Cu, 9% Zn
P0D-ECAL2	$210 \times 224 \times 30.4$	209.6	2.9				6.5% H, 40% C, 53.5% Pb
GArTPC1	$186 \times 206 \times 78$	251	0.016	$170 \times 196 \times 56$	256	0.010	100% Ar
FGD1	$186 \times 186 \times 30$	343	1.1	$175 \times 175 \times 29$	344	0.92	8% H, 88% C, 4% O
GArTPC2	$186 \times 206 \times 78$	387	0.016	$170 \times 196 \times 56$	256	0.010	100% Ar
FGD2	$186 \times 186 \times 30$	480	1.1	$175 \times 175 \times 29$	481	0.92	9% H, 50% C, 41% O
GArTPC3	$186 \times 206 \times 78$	524	0.016	$170 \times 196 \times 56$	256	0.010	100% Ar
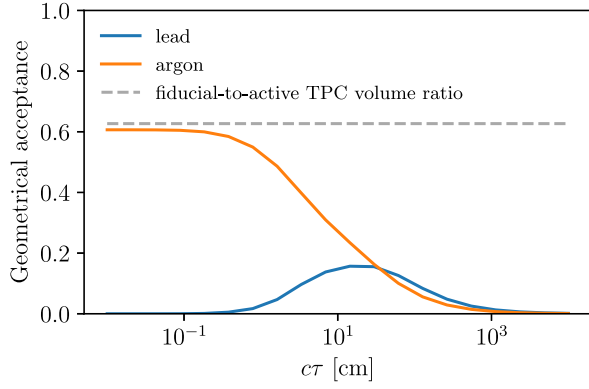
FIG. 4. The geometrical acceptance of ND280 as a function of the N proper lifetime. For the smallest lifetimes, only argon can pick up the decays.

the detector. In particular, we will focus on parameters such that the lifetimes of N are not much larger than $\mathcal{O}(10)$ m, as otherwise, these particles would not provide a good fit to MiniBooNE as well. If N propagates more than $\mathcal{O}(15)$ cm, it can be produced via upscattering in the dense material of the PØD, where the cross section is significantly enhanced due to the coherent scaling with proton number, Z^2. It would then decay into a visible e^+e^- pair inside one of the

TPCs. This particular signature is present in most of the interesting parameter space for the heavy-mediator case. When the lab-frame lifetime becomes shorter than $\mathcal{O}(15)$ cm, the heavy neutrinos produced in the PØD decay before entering the TPCs, and therefore the e^+e^- are rejected by the selection to avoid large neutrino-induced backgrounds. Nevertheless, the upscattering can happen inside the TPCs, where they would be visible. Despite the relatively small number of targets in the TPC fiducial volume, a handful of events is enough to constrain the model due to the absence of backgrounds. Such fast decays always happen in the light dark photon case ($m_{Z'} < m_N$), but it can also happen in regions of large $|V_N|^2$ values of the heavy dark-photon parameter space. Figure 4 shows the fraction of heavy neutrino decaying in one of the three TPCs as a function of the proper lifetime, for the case the upscattering happens in the lead or argon.

To estimate them, we developed a simplified detector simulation and implemented the analysis selection criteria on upscattering events generated by our own modified version of the DarkNews generator [76].

We expect differences in the reconstruction and selection efficiencies for upscattering with respect to the decay-in-flight signatures considered in Refs. [41,75]. Figure 5 shows the comparison between standard heavy neutrino
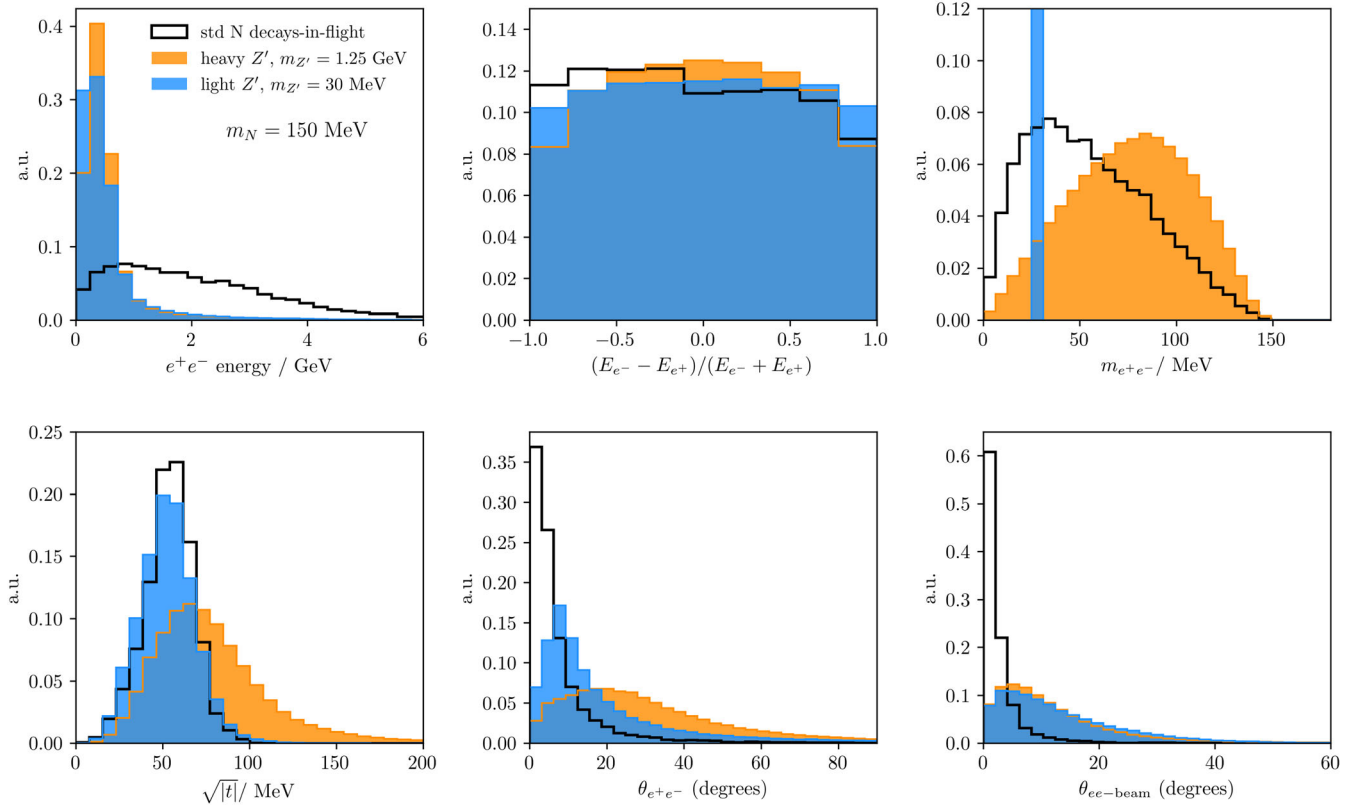

FIG. 5. Comparison between standard decay-in-flight $N \to \nu e^+e^-$ signatures of heavy neutral leptons produced at the target (solid black) and that of heavy neutrino decays initiated by coherent scattering in a light dark photon (filled blue) and heavy dark photon (filled orange) model. Histograms are area-normalized. For variable definitions, see Eq. (13).

signatures and our scattering-induced signatures for both the heavy and the light mediator case. The reconstruction efficiency depends on the kinematics of the heavy neutrino decay products, including the detector capability to separate the e^+e^- tracks as a function of their opening angle and their distance of closest proximity. While the selection efficiency relies on the following cuts [see Eq. (13)], which we implemented in our analysis. No smearing of the kinematics of the electron and positron has been applied at this stage.

$$E_{e^+e^-} \equiv E_{e^+} + E_{e^-} > 0.150 \text{ GeV}, \tag{13a}$$

$$m_{e^+e^-} \equiv \sqrt{(p_{e^+} + p_{e^-})^2} < 0.7 \text{ GeV}, \tag{13b}$$

$$|t| \equiv (E_{e^+e^-} - p_{e^+e^-}^z)^2 - |\vec{p}_{e^+e^-}^T|^2 < 0.03 \text{ GeV}^2, \tag{13c}$$

$$\cos\theta_{e^+e^-} \equiv \frac{\vec{p}_{e^+} \cdot \vec{p}_{e^-}}{|\vec{p}_{e^+}||\vec{p}_{e^-}|} > 0, \tag{13d}$$

$$\cos\theta_{ee-\text{beam}} \equiv p_{e^+e^-}^z / p_{e^+e^-} > 0.99. \tag{13e}$$

The efficiency for these cuts is of order 50% for our dark neutrino BPs. A more complete picture of the efficiency of these cuts across parameter space is given in Appendix D. Given that the efficiency in the original standard heavy neutrino analysis is of the order 10–15%, we applied an additional 10% efficiency factor to account for reconstruction effects in a conservative way.

We also perform a sensitivity study by projecting the status of this analysis by the end of the T2K data taking. The current analysis could be extended to about 4×10^{21} POT which have already been collected by the experiment. Moreover, ND280 is currently being upgraded to a new configuration [77]; the PØD is being replaced by two new GArTPCs and a Super-FGD module. A future search postupgrade, looking only at upscattering inside the argon, could be performed on the three TPCs, plus the two new TPCs, on a forecast of 16×10^{21} POT [78]. This conservative estimate neglects improvements to reconstruction and background rejection and a benefit of a tailored analysis for this model.

C. Analysis II: Photons in the FGD

The second analysis uses the photonlike control sample of the ν_eCCQE cross section measurement in the first FGD. Even though it focuses on a very different measurement, it can provide an important constraint for our model. The largest background for this analysis comes from photons that convert inside of the FGD and for which one of the two particles has not been reconstructed. In order to better measure this background, they look at a specific sideband, selecting e^+e^- in the FGD in the same way they select

single electrons or positrons for the primary measurement. The e^+e^- invariant mass is a helpful quantity for them to select real photons, and it can be used to constrain the dark neutrino signal in the case of a light mediator. The Z' is produced on shell and decays promptly to an e^+e^- pair, which, if reconstructed correctly, shows a peak in the invariant mass spectrum at $m_{ee} = m_{Z'}$. Figure 6 shows an example of the measured m_{ee} spectrum We implement the smearing of momenta and zenith angles using the matrices provided in Ref. [79]. We consider a flat 10% reconstruction and selection efficiency, which considers the 30% efficiency for the ν_eCCQE analysis, squared to account for the two leptons.

We also estimate the sensitivity of a projection of this analysis by expanding to two FGDs, with a larger dataset and including the SuperFGD after the upgrade.
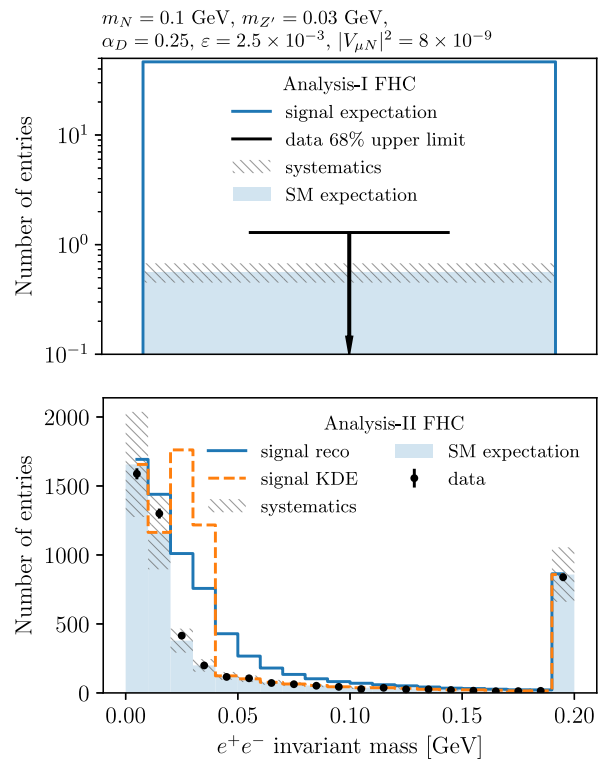

FIG. 6. The data we are using for Analysis-I (first panel) and Analysis-II (second panel) in neutrino mode (the plots for the antineutrino mode case are analogous). Analysis-I is a one-bin experiment, while Analysis-II is a search for a resonance on the e^+e^- invariant mass spectrum. The shaded blue region represents the expected SM background, while the black points show the observed data. No event was observed in Analysis-I. Therefore we display the upper limit at 68% C.L. The blue lines show the signal we expect to observe for the light Z' benchmark point but for a larger ε. The orange line illustrates the signal at the true level, smeared only by our KDE interpolation. It cannot be compared with the data but gives a sense of the effect of the experimental resolution in measuring e^+e^- pairs in the FGD.

IV. SIMULATION AND MONTE CARLO TECHNIQUES

Extended dark sectors like the ones considered in this work typically involve a large number of independent parameters. This observation poses a challenge from the phenomenology point of view. Performing inference in this ample parameter space requires predicting distributions of observables for all possible choices of parameters. It can be costly and quickly become unfeasible for more detailed simulations. For this reason, promptly predicting our signal is crucial to improving our coverage of the model parameter space. We apply existing statistical methods to interpolate model predictions across the parameter space, allowing the computation of the model prediction from a single simulated sample. Fast interpolations of physical predictions across the parameter space have been discussed in the context of event generators for colliders [80,81] as well as in fast generators of dark matter direct detection signatures [82]. These methods, however, rely on the parametrization of the prediction in terms of analytic functions. Our technique complements these methods by deriving a nonparametric estimate of the observables. A similar approach to the one discussed here has been proposed for treating nuisance parameters and systematic uncertainties in IceCube [83]. The IceCube scheme overcomes the curse of dimensionality of the production of many distinct Monte Carlo samples, sometimes described as the "multiple universes" approach. While the IceCube method derives a nonparametric estimate of the observables as a function of the nuisance parameters in the neighborhood of the central value, our method applies to the entire parameter space.

A. General idea

We now present the general idea behind our method. We can think of a model as a family of probability density functions (PDF) $p(x|\theta)$, where θ are the physical parameters of the theory over which we want to perform inference, like masses and couplings, while x are observables, like particle momenta. The model also predicts a normalization factor $\mathcal{N}(\theta)$; not just the observable distribution depends on the parameter, but also the total rate. Both θ and x are multidimensional, varying from several to $\mathcal{O}(10)$ dimensions.

Inference is performed by computing the expectation value $\mathrm{E}_\theta[T(x)]$ of a test statistic T for each value of θ. The typical approach proceeds as follows: (i) start from an initial definition of a multidimensional grid of a total of m points in the parameter space $\theta_{j=1,\ldots,m}$, (ii) run a simulation for each θ_j, i.e., draw n_j samples $x_{i=1,\ldots,n^j} \sim p(x|\theta_j)$, (iii) compute the expectation value for each θ_j as:

$$\mathrm{E}_{\theta_j}[T(x)] = \sum_{i=1}^{n^j} w_i^j T(x_i^j), \qquad (14)$$

where w_i^j are weights associated with the sampling, such as importance-sampling weights. Finally, (iv) one eventually

interpolates these values across the parameter space, in order to predict $\mathrm{E}_{\bar\theta}[T(x)]$ for a $\bar\theta$ which has not been simulated.

While the method above works, our procedure described next provides a more physical and complete way to interpolate the expectation. It allows us to rapidly compute multiple and more complex test statistics, like histograms, using a single set of samples, i.e., running only one simulation. We promote θ to a random variable by considering $p(x, \theta) = p(x|\theta)\mathcal{N}(\theta)q(\theta)$ where $q(\theta)$ is a prior over θ, and we sample $x_i, \theta_i \sim p(x, \theta)$ with weights w_i, for $i = 1, \ldots, n$. Using these samples, we obtain $\mathrm{E}_{\bar\theta}[T(x)]$ by interpolating across the parameter space using kernel density estimation,

$$\begin{aligned}
\mathrm{E}_{\bar\theta}[T(x)] &= \sum_i^n w_i T(x_i) \frac{w(\bar\theta, \theta_i)}{q(\theta_i)} \\
&= \sum_i w_i T(x_i) \frac{K(d(\bar\theta, \theta_i), \delta)}{q(\theta_i)}, \qquad (15)
\end{aligned}$$

where $K(d, \delta)$ is a kernel function, $d(\bar\theta, \theta_i)$ is a distance in parameter space, and δ is the bandwidth or smoothing parameter. By sampling over parameter space, we exploit the fact that neighbor parameters will produce similar observable distributions. Using some sampling techniques, like importance adaptive sampling or Markov Chain Monte Carlo, this method will guarantee to sample observables with significant contribution to any test statistic, i.e., large weights. However, the adaptation over the parameter space will also make to sample parameters where they result in larger weights. The function $q(\theta)$ allows us to control this effect and skew the distribution of samples towards our preferences. For example, if performing inference by conditioning on the posterior, e.g., when setting limits on a slice of the parameter space, fixing a subset of the total parameters, and varying the other ones, we might be in a region where there are no samples, as that slides contains a small probability with respect to the total model.

Finally, despite the power of this interpolation, it introduces some statistical uncertainty related to the finite sample size. Whenever a close formula for $w(\bar\theta, \theta_i)$ is present, it is more effective to use that. For this reason, we split the parameter space into θ^α, for which a KDE weight is computed, and θ^β, for which the weight is computed using an analytical formula. This procedure is illustrated in Fig. 7.

B. Application to dark neutrino sectors

We use this framework in the context of the dark neutrino model discussed in the previous sections. In this model, θ is a five-dimensional parameter space for the light mediator case and six-dimensional for the heavy case since $|V_N|$ is only a relevant variable for the latter. More precisely, $\theta^\alpha = \{m_N, m_{Z'}\}$, while $\theta^{\beta,\text{light}} = \{|V_{\mu N}|^2, \alpha_D, \varepsilon\}$, while
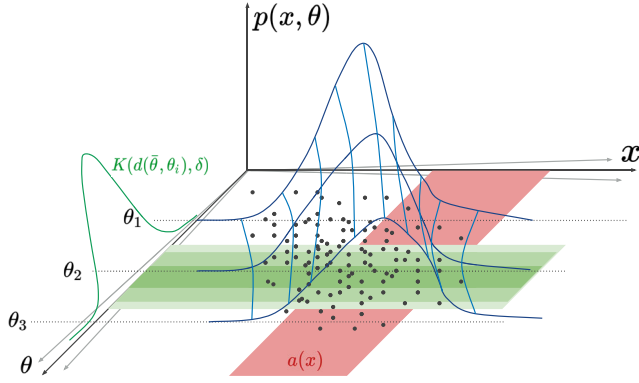
FIG. 7. An illustration of the construction of the posterior distribution using the kernel density estimator. In one set of axes, we show the model parameters, θ_i, and in another, we show the observable variables, x_i. In the vertical axis, we have the posterior distribution, $p(x, \theta)$. The black dots represent the samples thrown in the hyperspace of the model and phase space variables for which we calculate the posterior $p(x, \theta)$. The green and red planes are the support of the KDE, and the lines represent the posterior as interpolated by the KDE.

$\theta^{\beta,\text{heavy}} = \theta^{\beta,\text{light}} \cup \{|V_N|^2\}$. The differential cross section in the observable space $x = \{\Omega_N, \Omega_{ee}\}$ is our $p(x|\theta)$, as

$$p(x|\theta) = \frac{d\sigma}{d\Omega} = \frac{d\sigma}{d\Omega_N}(\Omega_N|\theta) \frac{1}{\Gamma(\theta)} \frac{d\Gamma}{d\Omega_{ee}}(\Omega_{ee}|\Omega_N, \theta), \quad (16)$$

where Ω_N describes the kinematics of the heavy neutrino and the nuclear recoil, while Ω_{ee} describes the kinematics of the e^+e^- pair and the final state neutrino. However, since we are not interested in the incoming neutrino energy, sampled according to the flux and the degrees of freedom of the recoil and the final state neutrino, we implicitly integrate over those variables. Here, Γ is the total decay width of the heavy neutrino and is given by

$$\Gamma(\theta) = \int d\Omega_{ee} \frac{d\Gamma}{d\Omega_{ee}}(\Omega_{ee}|\Omega_N, \theta), \quad (17)$$

and can be computed analytically using Eqs. (9) and (10) for the heavy and light scenarios, respectively. Finally, an essential parameter for the simulation of the events in the detector is the heavy neutrino lifetime in its rest frame $\tau^0(\theta) = \hbar/\Gamma(\theta)$.

C. Monte Carlo event generator

We implemented the physics matrix elements in a Monte Carlo event generator, and we sample events using the Vegas Monte Carlo algorithm [84,85] with its Python implementation [86].

In a typical simulation, we would sample the following integral

$$\sigma(\theta) = \int d\Omega_N \frac{d\sigma}{d\Omega_N}(\Omega_N|\theta) \int d\Omega_{ee} \frac{1}{\Gamma(\theta)} \frac{d\Gamma}{d\Omega_{ee}}(\Omega_{ee}|\Omega_N), \quad (18)$$

that we now extend to

$$\int d\theta^\alpha q(\theta^\alpha) \int d\Omega_N \frac{d\sigma}{d\Omega_N}(\Omega_N|\theta)$$
$$\times \int d\Omega_{ee} \frac{1}{\Gamma(\theta)} \frac{d\Gamma}{d\Omega_{ee}}(\Omega_{ee}|\Omega_N), \quad (19)$$

where $\theta^\alpha = \{m_N, m_{Z'}\}$. We use $q(\theta^\alpha)^{\text{light}} = m_{Z'}^2/m_N^{3.5}$ and $q(\theta^\alpha)^{\text{heavy}} = m_{Z'}^8/m_N^5$, designed to provide samples distributed as uniformly as possible in the $\{m_N, m_{Z'}\}$ plane.

We employ the total number of selected events in a single-bin analysis as test statistics,

$$\mu(\theta) = \sum_k n_t^k \times \text{POT} d\Omega_N \frac{d\sigma}{d\Omega_N}(\Omega_N|\theta)$$
$$\times \frac{1}{\Gamma(\theta)} d\Omega_{ee}(\Omega_{ee}|\Omega_N)\epsilon(\Omega_{ee})a(\Omega_N, \Gamma(\theta)), \quad (20)$$

where the cross section has been multiplied by n_t^k, the number of targets for each material, indexed by k, and by the collected beam exposure in terms of protons on target. We also folded in the selection efficiency $\epsilon(\Omega_{ee})$ and detector acceptance $a(\Omega_N, \Gamma(\theta))$, which depends on the kinematics as well as on the lifetime of the heavy neutrino. Both functions can be computed as multidimensional cuts on the observables.

By introducing weights for parameters θ^α and θ^β such that,

$$\text{E}_{\bar{\theta}}[T] = \int d\Omega_N d\Omega_{ee} T(x) \frac{d\sigma}{d\Omega_N}(\Omega_N|\bar{\theta}) \frac{1}{\Gamma(\bar{\theta})}(\Omega_{ee}|\Omega_N)$$
$$= \sum_i^n w_i T(x_i) w_i^{\text{KDE}}(\bar{\theta}^\alpha, \theta_i^\alpha) w_i^\sigma(\bar{\theta}^\beta, \theta_i^\beta), \quad (21)$$

and by defining $\epsilon(\Omega_{ee,i}) = w_i^\epsilon(\Omega_{ee,i})$, $a(\Omega_{N,i}, \Gamma(\theta_i)) = w_i^{\tau_0}(\theta_i, \Omega_{N,i})$, and $n_t^{k_i} \times \text{POT} = w_i^{n_t, \text{POT}}$ if event i is generated on material k_i, we can rewrite Eq. (20) as a product of weights,

$$\mu(\theta) \simeq \sum_i^n w_i w_i^{\text{KDE}}(\bar{\theta}^\alpha, \theta_i^\alpha) w_i^\sigma(\bar{\theta}^\beta, \theta_i^\beta)$$
$$\times w_i^\epsilon(\Omega_{ee,i}) w_i^{\tau_0}(\theta_i, \Omega_{N,i}) w_i^{n_t, \text{POT}}. \quad (22)$$

For simplicity, we will discuss the method by writing the expectation for a single-bin analysis, which is the case for Analysis-I. Binned analyses, like Analysis-II, represent a trivial extension.

D. Multidimensional reweighting scheme

We now detail the various weights appearing in Eq. (22). Figure 8 summarizes the different aspects of the reweighting scheme.

a. Cross section KDE weights: The cross-section has a nontrivial dependence on $\theta^\alpha = \{m_N, m_{Z'}\}$. We define the KDE weight as

$$w_i^{\mathrm{KDE}}(\bar{\theta}^\alpha, \theta_i^\alpha) = K(d(\bar{\theta}^\alpha, \theta_i^\alpha), \delta)/q(\theta_i^\alpha). \quad (23)$$

We studied the accuracy of different kernels, distance functions, and smoothing parameters by comparing the interpolation on a benchmark grid with a dedicated, high-statistics sample for different values of θ. In the rest of the work, we used the Epanechnikov kernel, a logarithmic distance, and $\delta = 0.005$ along the direction of both parameters, in an uncorrelated way.

b. Cross section analytical weights. The upscattering cross section is proportional to $|V_{\mu N}|^2 \alpha_D (e\epsilon)^2$, as seen by squaring the amplitude in Eq. (4). We implement a trivial scaling along with the parameters θ^β allowing a quick reweight of the events in this parameter space. In this case, we define

$$w_i^\sigma(\bar{\theta}^\beta, \theta_i^\beta) = \frac{|V_{\mu N}|^2 \alpha_D (e\epsilon)^2}{(|V_{\mu N}|^2 \alpha_D (e\epsilon)^2)_i}, \quad (24)$$

where the parameters $\bar{\theta}^\beta$ are fixed in the entire simulation and so independent of i, but could, in principle, be varied as well.

c. Reconstruction and selection efficiency weights: We implemented $\epsilon(\Omega_{ee})$ as a function that is 0 for the events which are not selected and $\bar{\epsilon}$ for the events that are selected. For Analysis-I, this weight is $\bar{\epsilon} = 10\%$ and the selection

follows Eq. (13). For Analysis-II, this weight is a flat $\bar{\epsilon} = 10\%$ for every event.

d. Lifetime reweighting: This weight applies only to the heavy case, where lifetimes span multiple orders of magnitude, while the light case always leads to a prompt decay ($c\tau^0 \leq 0.1$ cm). In this latter case, we simulate interactions directly in the fiducial volume. The easiest way to compute the acceptance for different lifetimes is to sample a number from the exponential distribution with a scale parameter equal to the N lab-frame lifetime, propagate N to the detector, and accept or reject the event if the decay point happens within the TPC fiducial volume. However, this method has an important drawback as it produces small effective sample sizes, especially at short lifetimes, where most interactions from the PØD will not make it to the detector. To avoid this issue, we instead account for the geometrical acceptance by multiplying by a lifetime weight, which is equal to the integral of the trajectory within the TPC weighted by the exponential distribution. The trajectory of the heavy neutrino in the lab frame in the event i enters and exits each of the three different TPC at points (a_i^j, b_i^j), where $j = 0, 1, 2$ is the TPC index. If the heavy neutrino never enters a given TPC, we can take both numbers as infinity. For each event, we can compute $(\beta\gamma)_i = p_i/m_N$, and given a value of the lifetime in the proper frame $c\tau_0$, the lifetime weight is computed as

$$
\begin{aligned}
w_i^{\tau_0}(\theta, \Omega_N) &= \sum_j \int_{a_i^j}^{b_i^j} \frac{ds}{(\beta\gamma)_i c\tau_0} e^{-s/(\beta\gamma)_i c\tau_0} \\
&= \sum_j (e^{-a_i^j/(\beta\gamma)_i c\tau_0} - e^{-b_i^j/(\beta\gamma)_i c\tau_0}). \quad (25)
\end{aligned}
$$

e. POT and number of targets weights: Reweighting for the POT and number of targets is trivial, as it depends only on these multiplicative factors

$$w_i^{\mathrm{n_t, POT}} = \mathrm{n_t} \times \mathrm{POT}, \quad (26)$$

and can be computed on the fly to change the beam exposure and target material and mass easily.

E. Likelihood evaluation

We compute a Poisson likelihood of the observed data (N_{obs}) given the expectation, summing up the expected background [b and the signal ($\mu(\theta)$] across the parameter space. We account for systematic uncertainties by using the *effective likelihood* framework [87], which provides an analytic formula to marginalize over systematic uncertainties,

$$\mathcal{L}(\theta) = \mathcal{L}^{\mathrm{eff}}(\theta|N_{\mathrm{obs}}, b + \mu(\theta), \sigma^2(b, \mu(\theta))), \quad (27)$$
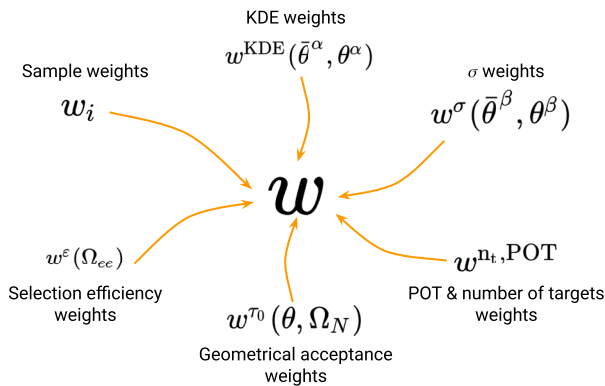
where

FIG. 8. In order to quickly compute the model prediction across the parameter space, we implement a set of weights that accounts for the simulation of the cross section, the kinematics, and detector effects. By taking the product of all these weights, we obtain a single weight that can be used to compute the final model prediction for any observable.
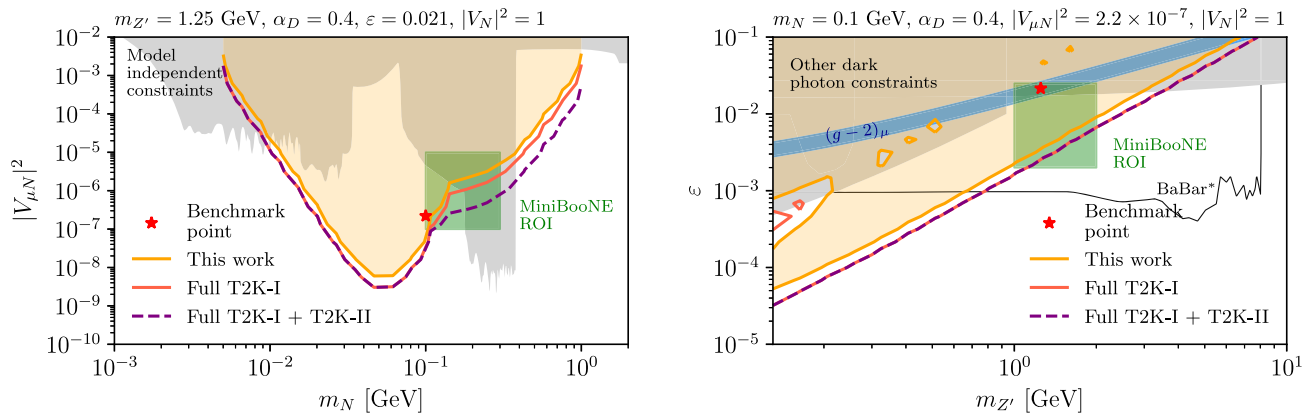
FIG. 9. Limits on the dark neutrino model for a scenario with a heavy dark photon, where $c\tau_N^0$ is typically greater than centimeters. The red star corresponds to benchmark point B. On the left, we show the limits on $|V_{\mu N}|^2$ as a function of m_N and on the right on ϵ as a function of $m_{Z'}$, choosing the remaining parameters according to benchmark (B). The orange-shaded region is excluded by our analysis using the available T2K data. The red line shows the sensitivity estimate for the analysis based on the T2K data collected during the first phase, before the ND280 upgrade, while the purple line includes all the data that will be collected by the end of the life of the experiment. The MiniBooNE region of interest (ROI) is shown as a large green area surrounding the benchmark point in Ref. [29].

$$\sigma^2(b, \mu(\theta)) = \sum_i w_i^2 + (b + \mu(\theta))^2 * \eta^2, \qquad (28)$$

accounts for systematic uncertainties. The first addend accounts for the finite sample size, while the second includes the analysis systematics (e.g., flux and cross section), using the fractional systematic uncertainties published with the analysis, which are typically close to a flat 20%. This formula can be easily extended to a multibin analysis by taking the product of the likelihood for each bin. When combining different analyses, like the TPC search and the FGD sideband, we simply sum the likelihood together. When computing projections, we scale the signal and the background proportional to the number of targets and the POT, and we assume $N_{\text{obs}} = \text{int}(b + \mu(\theta))$, where int() is just approximating to an integer number.

V. RESULTS AND DISCUSSION

Given that neither analysis observed any excess of events with respect to the background prediction, they can be used to constrain the parameter space of the dark neutrino model. We show our limits for two particular projections in parameter space; the $m_N - |U_{\mu N}|^2$ plane, describing the heavy neutrino properties, and the $m_{Z'} - \epsilon$ plane, describing the dark photon properties. We compute the likelihood on the plane by summing up the negative log-likelihoods of the relevant analyses, subtracting the minimum, and tracing the contour at constant 2.3, which produces regions of exclusion at 90% C.L. These contours are shown in Fig. 9 for the heavy case and Fig. 10 for the light case. We discuss the shapes of the excluded regions and the individual contributions of the different parts of the detector to the analyses in Appendix E.
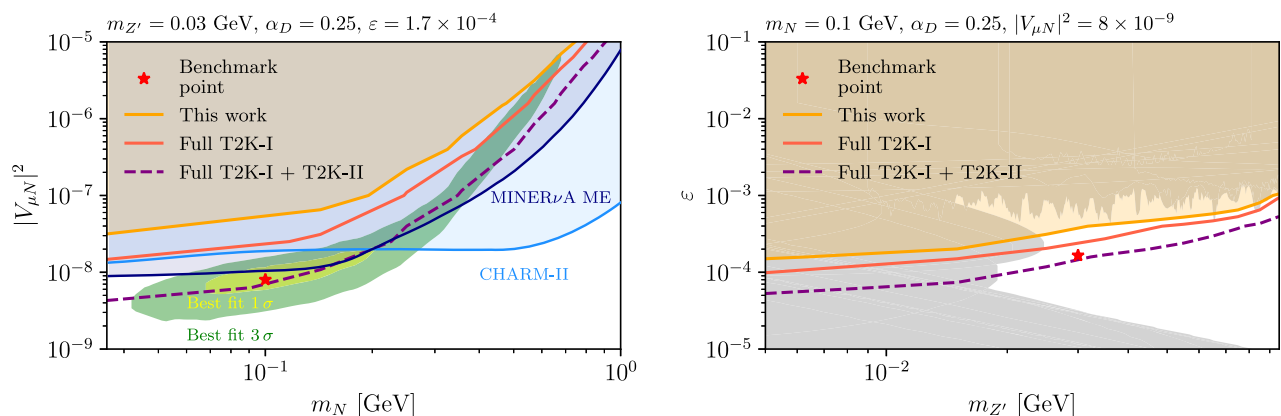

FIG. 10. Limits on the dark neutrino model for a scenario with a light dark photon, where N decays are prompt. The red star corresponds to Benchmark point A. On the left we show the limits on $|V_{\mu N}|^2$ as a function of m_N for $m_{Z'} = 30$ MeV and on the right on ϵ as a function of $m_{Z'}$ for $m_N = 100$ MeV. The orange shaded region, the red and the purple lines have the same meaning as in Figure 9. We show the allowed region from Ref. [25] on the left. Limits on the light and visible dark photon have been obtained from Ref. [88].

Given that no fit of this model to the MiniBooNE data has been performed for the heavy case, we consider a region of interest around the benchmark point, while for the light case, we consider the best-fit region from [25]. For the heavy case, we also determine model-independent bounds on the proper decay length $c\tau^0$ as a function of the upscattering rate after selection but before requiring that the decay happened in the TPC, as shown in Fig. 11. We derive approximated functional forms for the bounds, which could easily be applied to other models predicting similar signatures. More details about the derivation and the validity of these functional forms are outlined in Appendix F.

As previously discussed in Ref. [41,75,89], the GAr time-projection chambers of the T2K near detector, ND280, provide a powerful probe of long-lived particles. Visible decays inside the low-density volume of the TPCs,

where neutrino-induced backgrounds are negligible, are clearly identified. In this work, we showed that the combination of the high-density material in the PØD detector with the low-density TPCs downstream is even more powerful. The former enhances the production of new particles in neutrino-nucleus scattering due to the large mass of lead. At the same time, the latter provides a desirable volume to search for charged final states. In addition, the magnetic field allows for improved identification of $\ell^+\ell^-$ pairs even at the smallest opening angles and energies.

In Analysis-I, we use an existing, background-free search for e^+e^- pairs inside the volume of the TPCs [41] to set limits on a short-lived heavy neutrino N, produced in neutrino-nucleus upscattering upstream in the detector. Having observed zero events, T2K data strongly constrains these particles as explanations of the MiniBooNE excess but does not entirely rule them out for
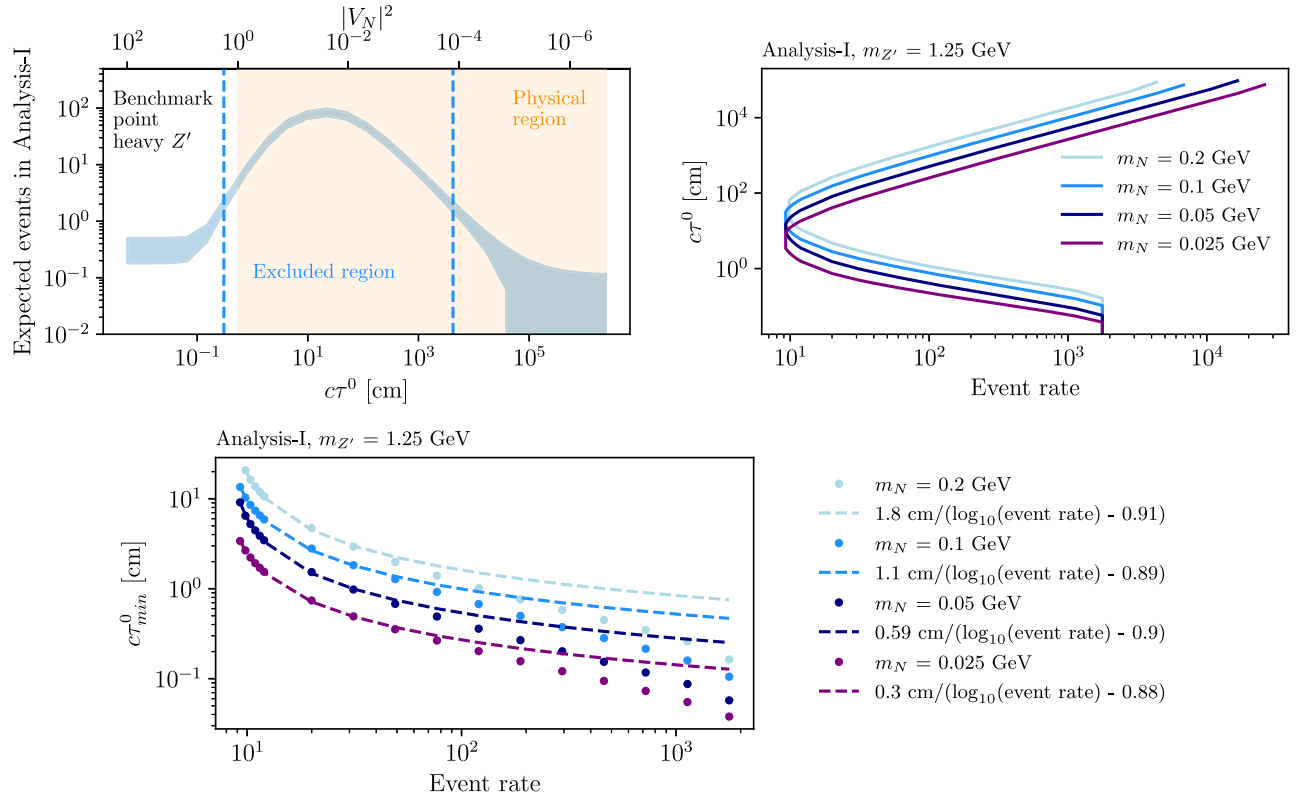

FIG. 11. Top left: The expected number of events in Analysis-I after the full selection as a function of the proper decay length of N, $c\tau^0$, for the heavy Z' benchmark point. The blue band represents the uncertainty as obtained using Eq. (28). Once we fix all the other parameters in the model, there is a bijection between $c\tau^0$ and $|V_N|^2$, as shown on the top x-axis. However, $|V_N|^2$ is physical only in the region allowed by Eq. (11), shaded in orange. Besides very long-lived heavy neutrinos, we exclude most of the physical region. The minimum decay length we exclude is 0.3 cm, although our model cannot generate such short lifetimes for this combination of parameters. Top right: The limits on the N decay length as a function of the total number of events after signal selection but before requiring that N decays in the TPC volume. The region enclosed by the solid lines is excluded at 90% C.L. At first-order approximation, different $m_{Z'}$ values result in different event rates without affecting the kinematics of the decay, while different m_N values result in different production thresholds and lab-frame lifetimes, explaining the difference between the different curves. Bottom: Zoom of the previous plot in the lower-left part of the graph. The dots represent the curves obtained with the analysis, while the dashed lines are the best-fit values obtained with a simplified model. Although the functional form cannot perfectly fit all the points, it describes the curves reasonably accurately, providing model-independent constraints which can be applied to models predicting a similar phenomenology.

sufficiently short lifetimes, $c\tau_N^0/m_N \lesssim 2$ cm/GeV. In Analysis-II, we explored the measurement of single-photon events in the Fine-Grained Detector [71] to set limits on the parameter space with the shortest lifetimes of N. Due to the larger backgrounds and lower event rate, this analysis is not as sensitive as Analysis-I and therefore does not exclude allowed regions of parameter space. These limits can be significantly improved using all T2K-I run data taken up to 2021, as well as with future data that will be taken with the upgraded ND280 detector and the more intense neutrino beam at J-PARC.

We note that we have not exhausted the list of models, having not covered cases with small mass splittings between N and the daughter neutrinos, scalar mediator models, and other $2 \rightarrow 3$ scattering signatures involving the emission of on-shell dark bosons [32–37,90]. We expect different signal selection efficiencies for these models, especially those that better fit the MiniBooNE angular spectrum. We encourage the T2K collaboration to pursue a dedicated search for all such upscattering signatures, including the one discussed in this paper, leveraging the full power of their detector simulation. In particular, we expect that a complete reconstruction simulation by the Collaboration overcomes the simplifying assumption in this work of energy-independent signal efficiencies. We expect our method to explore the parameter space to prove even more beneficial in the experimental setup, where it can reduce the computational cost of full simulations, reducing the sampling in un-interesting regions of parameter space. In addition, further public data on the reconstruction efficiencies as a function of physical observables, like energies and angles, rather than model parameters, would be incredibly beneficial to the phenomenology community.

VI. CONCLUSIONS

We have proposed a method to cover multidimensional model parameter spaces efficiently. We used importance sampling to sample model parameters and physical observables to find model predictions for the number of events in a particular experimental setup. Because both types of parameters are sampled intelligently, the construction of the likelihood is fast, flexible, and efficient. We also built a kernel density estimator to interpolate the likelihood inside the sampled region of parameter space. The KDE gives us greater flexibility in hypothesis testing and offers a valuable product that can be distributed to other users interested in the model likelihood.

The interpolation is achieved with a single batch of simulated events and can provide limits in arbitrary slices of the parameter space. We illustrated the power of this tool by setting limits on the parameter space of a model of great interest in neutrino physics, as shown in the four planes of Figs. 9 and 10. In the context of particle physics experiments, this method will constitute a useful tool for experimental collaborations and phenomenologists to explore more complex physical models. In neutrino physics, with the latest progress in building new phenomenology-friendly Monte Carlo generators for new physics processes [91–93], improving our ability to cover the large-dimensional parameter spaces of dark sectors will be all the more relevant.

More technical studies are needed for a quantitative comparison between the efficiency of our method and standard grid sampling. Clearly, the choice of a model-parameter-space sampling method will strongly depend on the case considered. Nevertheless, one can expect that adaptive sampling and KDEs will be more advantageous for complex, smoothly varying, and large-dimensional parameter spaces. Other future directions include applying our methodology to searches for new physics outside the context of short-baseline anomalies. Among models of interest are higgsed low-scale $U(1)$ symmetries, coannihilating dark matter models, light axionlike particles with multiple interactions, and flavor-rich models of heavy neutrinos. In these theories, decay-in-flight signatures of multiple light dark particles are ubiquitous and require more efficient exploration methods to be thoroughly tested. We believe our method can increase the theory reach of experiments like the Large Hadron Collider, the Short-Baseline Neutrino program at Fermilab [94,95], atmospheric neutrino experiments like IceCube and KM3NET [96], as well as future high-intensity long-baseline experiments like DUNE [97] and Hyper-Kamiokande [98].

ACKNOWLEDGMENTS

APPENDIX A: UV COMPLETIONS

Possible UV completions of Eq. (1) have been discussed in Refs. [26,29,31]. The general idea is to consider new fermions ν_D charged under the new gauge symmetry, which mix with SM neutrinos upon symmetry breaking. Two main categories can be identified, depending on the pattern of the $U(1)_D$ breaking. Schematically, they make use of the following operators,

$$(\text{I}) : (\bar{L}\tilde{H}_D)\nu_D, \tag{A1}$$

$$(\text{II}) : (\bar{L}\tilde{H})(\Phi\nu_D). \tag{A2}$$

The first route requires new $SU(2)_L$ scalar doublets, H_D, also charged under the dark symmetry. The mixing between SM neutrinos and the dark leptons ν_D is then generated by the expectation value of H_D, which breaks the $U(1)_D$, together with the SM Higgs, also the electroweak symmetry. The second method considers instead an SM-singlet dark scalar Φ, whose expectation value breaks only the dark symmetry. The main idea is illustrated by the dimension-six operator in Eq. (A2), which induces a mixing term between ν_D and the SM neutrinos after symmetry breaking. This effective operator can be easily generated by the exchange of a singlet (sterile) neutrino ν_s, which serves as a bridge between the SM and the dark sector via $(\bar{L}\tilde{H})\nu_s$ and $\overline{\nu_s}(\nu_D\Phi)$.

Note that in both cases, the dark photon gets a mass from breaking $U(1)_D$, while the masses of the neutral leptons and the additional scalar degrees of freedom will depend on the specifics of the model. The interplay between the expectation values of the new scalars, Yukawa couplings, and new arbitrary Majorana mass terms will determine the coupling vertices V_{ij} in Eq. (2) and should also generate the correct value for the light neutrino masses. In particular, both types of model are flexible enough for N to be a (pseudo-)Dirac or Majorana particle. However, pseudo-Dirac states are preferred to generate small neutrino masses due to the approximate conservation of the lepton number. Any additional fermion in the model can be heavier than a few GeV, where their interactions with the SM would be poorly constrained at the values of neutrino mixing we consider.

APPENDIX B: ANALYTICAL APPROXIMATION FOR UPSCATTERING CROSS SECTION

For convenience, a crude approximation for the upscattering cross sections above $E_\nu > 1$ GeV is given below. These have been obtained assuming a box function for the coherent and dipole form factors with cutoffs around the QCD scale and vector mass, respectively. For upscattering on nuclei,

$$\sigma_{\text{coh}}^{\nu_\alpha \to N} \simeq \frac{|V_{\alpha h}|^2 (Ze\epsilon)^2 \alpha_D}{4E_\nu^2 m_{Z'}^4} [2(M^4 + s^2) - sM^2(x_A^2 + 4)], \tag{B1}$$

where Z is the atomic number of the nucleus with mass M, $x_A = 2\Lambda_{\text{QCD}}/A^{1/3}$ with Λ_{QCD} from 100 MeV to 200 MeV, and A the atomic mass number. The dependence of the total cross section on Λ_{QCD} is stronger for lower energies.

APPENDIX C: DETECTOR DESCRIPTION

We simulate the three subdetectors of ND280: the PØD, the two FGDs, and the three GAr TPCs. In Table II, we report the sizes of the active volume, where upscaterring occurs, and the fiducial volume, where e^+e^- pairs are detected. We also report total active and fiducial mass, as well as the material composition in mass. We account for gaps between the detector volumes and report the z coordinate along the beam axis, where the active or fiducial volume begins.

Figure 12 shows the 2D distribution density of upscattering vertices along the z and x axes for the heavy Z' case, using our benchmark point in parameter space. The three different sections of the PØD, the three different TPCs, and the two FGDs are clearly distinguishable, together with the gaps between volumes.
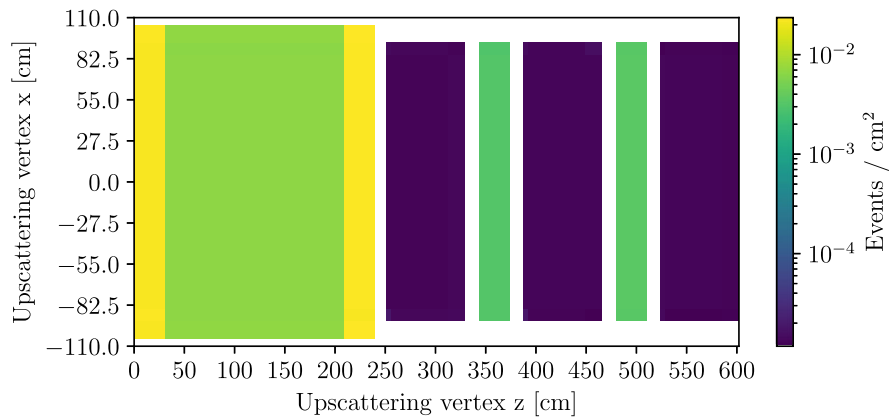

FIG. 12. Distribution of the upscattering vertex across the z-x plane, for the heavy Z' case, using our benchmark point. The color bar shows the number of upscattering events per bin.
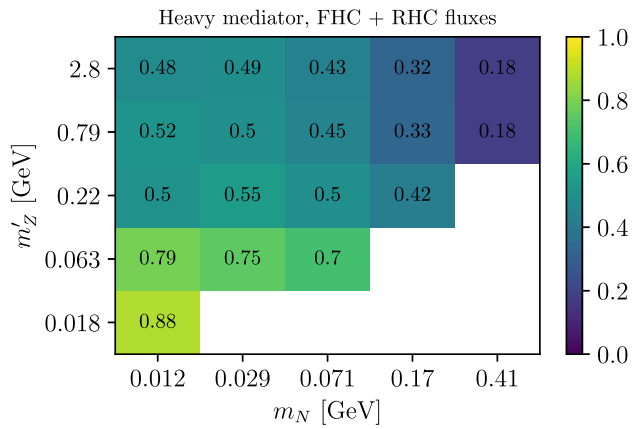
FIG. 13. Selection efficiency resulting from the kinematical cuts in Eq. (13) in Analysis-I for the heavy mediator case, as a function of the parameter space $m_{Z'}-m_N$.

APPENDIX D: KINEMATICAL SELECTION EFFICIENCY IN ANALYSIS-I

Figure 13 shows a map of the efficiency for the kinematical cuts in Eq. (13) as a function of m_N and $m_{Z'}$, for the heavy Z' case in Analysis-I. The distributions of the kinematical variables and, therefore, the selection efficiency depend only on these two parameters because all the others affect only the total upscattering rate or the lifetime. The efficiency is about 50% for our benchmark point, and it is independent of $m_{Z'}$ for large values of $m_{Z'}$ when $m_{Z'}$ only determines the total rate, similar to the W mass in the muon decay. However, it grows to almost 90% at lower values of $m_{Z'}$, while it decreases to $\sim 20\%$ at larger m_N. While automatically considered when scanning the parameter space through our procedure, this variation does not significantly impact the final result.
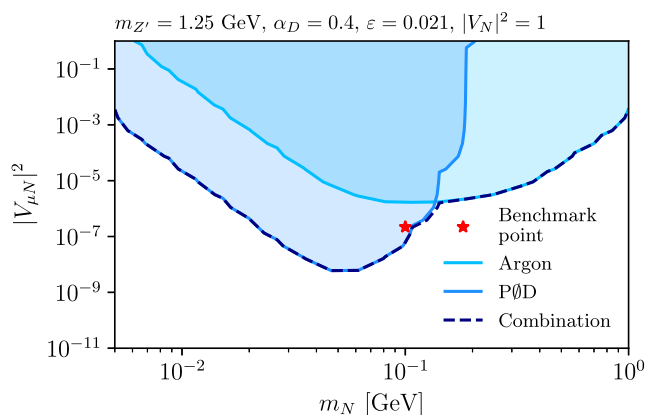
APPENDIX E: COMPLEMENTARITY OF THE ND280 SUBDETECTORS

1. Upscattering in the P∅D and in the GArTPCs

In the heavy mediator case, both the scattering in the P∅D and the argon contribute significantly to the constraints, however, in different regions of the parameter space. Scattering in the gaseous argon is rare because of the low density, but it is the most powerful component in constraining the shortest lifetimes since it is where the fiducial volume of the analysis is contained. Figure 14 shows our constraints, as in Fig. 9, splitting the limits into the contribution from the GArTPCs and from the P∅D. Between $m_N = 0.1$ GeV and $m_N = 0.2$ GeV, the model is very short lived, and all heavy neutrinos produced in the P∅D decay before reaching the TPCs. This region is constrained only by prompt decays of heavy neutrinos produced inside the argon and is, therefore, less constrained. In the right plot, we show the dark photon parameter space, where, despite the larger upscattering rate at smaller Z' masses, the model cannot be constrained by the P∅D events due to the short lifetimes. However, in several regions of parameter space, the model predicts a significant number of events in the argon, which allows for a robust exclusion of the largest values of ε.

2. Upscattering in the FGD and the GArTPCs

In the light mediator case, we combine Analysis-I and Analysis-II, considering upscattering happening in the GArTPCs, for the first case, and in the FGDs, in the second case. The two analyses contribute similarly to the limit, as shown in Fig. 15.
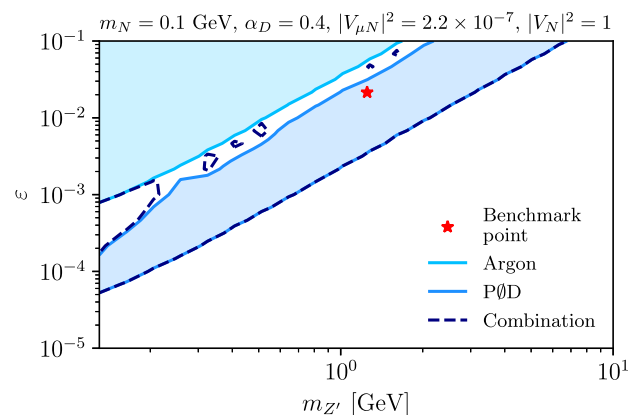

FIG. 14. Limits for the heavy Z' case, as shown in Fig. 9, splitting into the two regions where upscattering happens in the GArTPCs or the P∅D.
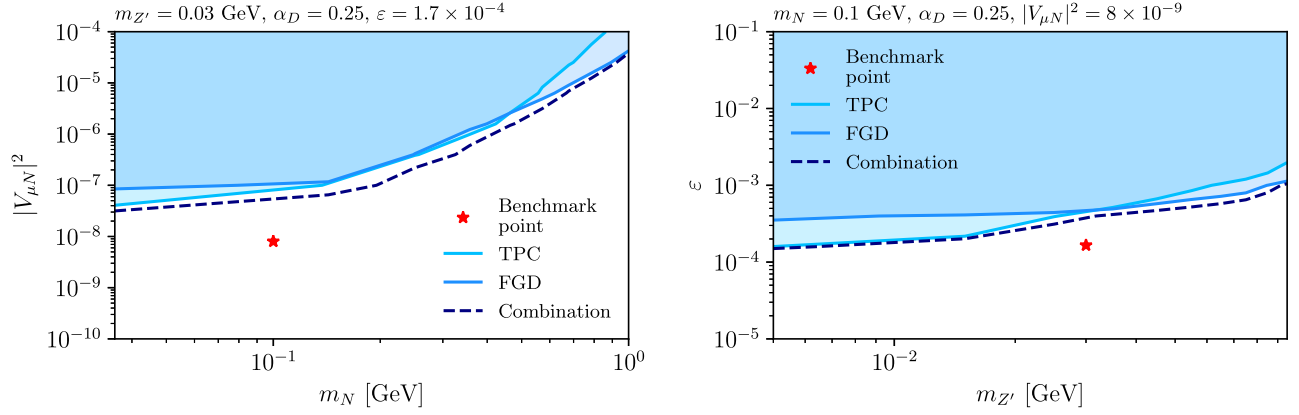
FIG. 15. Limits for the heavy Z' case, as shown in Fig. 10, splitting into the two regions where upscattering happens in the FGD or the GArTPCs.

APPENDIX F: FUNCTIOAL FORMS OF THE MODEL-INDEPENDENT CONSTRAINTS

The bottom plot of Fig. 11 shows two limits; the dots represent the limits derived through the full analysis, corresponding to lines of the top right figure, while the dashed lines show best-fit curves using the well-motivated functional form,

$$ c\tau^0_{\min} = \frac{\alpha[\text{cm}]}{\log_{10}(\text{event rate}) - \beta}, \tag{F1} $$

where α and β are free parameters. It can be derived by assuming a pointlike source of heavy neutrinos decaying in a detector of size d at a distance z and assuming that the process happens in one dimension only. The number of decays in the detector is given by event rate $\times$ $(e^{-z/\beta\gamma c\tau^0} - e^{-(z+d)/\beta\gamma c\tau^0})$. We are considering the region of the smallest lifetimes we exclude, so $c\tau^0 \ll z$, and the second term of the exponential is negligible. Finally, the limit is set when the event rate is equal to a constant value C, so the equation defining the limit can be rewritten as

$C = e^{-z/\beta\gamma c\tau^0}$, with $\alpha = z/\beta\gamma$ and $\beta = f(C)$ as free parameters, with f a certain functional form. It can now be inverted to obtain Eq. (F1).

The best fit to the different curves accurately describes the points at small-event rates, underestimating the limit at smaller lifetimes. As expected, the values of β are roughly constant, while α varies with m_N, roughly in a proportional way. The reason is that the momentum available in the heavy neutrino production is constant, so $m_N\beta\gamma$ is roughly constant. Combining these two observations, we deduce that we exclude

$$ \frac{c\tau^0}{m_N} \gtrsim \frac{1.1 \text{ cm}}{0.1 \text{ GeV}} \frac{1}{\log_{10}(\text{event rate}) - 0.9}. \tag{F2} $$

While this functional form is limited and cannot replace the complete study, it provides a model-independent benchmark, which can easily show the size of the constraints resulting from this analysis for an arbitrary model without needing a complete analysis.

[1] P. Minkowski, $\mu \to e\gamma$ at a rate of one out of 10^9 muon decays?, Phys. Lett. **67B**, 421 (1977); R. N. Mohapatra and G. Senjanovic, Neutrino Mass and Spontaneous Parity Violation, Phys. Rev. Lett. **44**, 912 (1980); M. Gell-Mann, P. Ramond, and R. Slansky, Complex spinors and unified theories, Conf. Proc. C **790927**, 315 (1979), arXiv:1306 .4669; T. Yanagida, Horizontal symmetry and masses of neutrinos, Conf. Proc. C **7902131**, 95 (1979); G. Lazarides, Q. Shafi, and C. Wetterich, Proton lifetime and fermion masses in an SO(10) model, Nucl. Phys. **B181**, 287 (1981); R. N. Mohapatra and G. Senjanovic, Neutrino masses and mixings in gauge models with spontaneous parity violation, Phys. Rev. D

23, 165 (1981); J. Schechter and J. W. F. Valle, Neutrino masses in SU(2) × U(1) theories, Phys. Rev. D **22**, 2227 (1980); T. P. Cheng and L.-F. Li, Neutrino masses, mixings and oscillations in SU(2) × U(1) models of electroweak interactions, Phys. Rev. D **22**, 2860 (1980); R. Foot, H. Lew, X. G. He, and G. C. Joshi, Seesaw neutrino masses induced by a triplet of leptons, Z. Phys. C **44**, 441 (1989).

[2] A. M. Abdullahi et al., The present and future status of heavy neutral leptons, arXiv:2203.08039.

[3] A. A. Aguilar-Arevalo et al. (MiniBooNE Collaboration), A Search for Electron Neutrino Appearance at the $\Delta m^2 \sim 1 \text{ eV}^2$ Scale, Phys. Rev. Lett. **98**, 231801 (2007).

[4] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations, Phys. Rev. Lett. **105,** 181801 (2010).

[5] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Improved Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations in the Mini-BooNE Experiment, Phys. Rev. Lett. **110,** 161801 (2013).

[6] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Significant Excess of ElectronLike Events in the Mini-BooNE Short-Baseline Neutrino Experiment, Phys. Rev. Lett. **121,** 221801 (2018).

[7] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Updated MiniBooNE neutrino oscillation results with increased data and new background studies, Phys. Rev. D **103,** 052002 (2021).

[8] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), MiniBooNE and MicroBooNE Joint Fit to a $3+1$ Sterile Neutrino Scenario, Phys. Rev. Lett. **129,** 201801 (2022).

[9] C. Athanassopoulos *et al.* (LSND Collaboration), Evidence for neutrino oscillations from muon decay at rest, Phys. Rev. C **54,** 2685 (1996).

[10] C. Athanassopoulos *et al.* (LSND Collaboration), Evidence for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations from the LSND Experiment at LAMPF, Phys. Rev. Lett. **77,** 3082 (1996).

[11] C. Athanassopoulos *et al.* (LSND Collaboration), Evidence for $\nu_\mu \to \nu_e$ Neutrino Oscillations from LSND, Phys. Rev. Lett. **81,** 1774 (1998).

[12] C. Athanassopoulos *et al.* (LSND Collaboration), Evidence for $\nu_\mu \to \nu_e$ oscillations from pion decay in flight neutrinos, Phys. Rev. C **58,** 2489 (1998).

[13] A. Aguilar-Arevalo *et al.* (LSND Collaboration), Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam, Phys. Rev. D **64,** 112007 (2001).

[14] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, Astron. Astrophys. **641,** A6 (2020); **652,** C4(E) (2021).

[15] M. A. Acero *et al.*, White paper on light sterile neutrino searches and related phenomenology, arXiv:2203.07323.

[16] S. N. Gninenko, The MiniBooNE Anomaly and Heavy Neutrino Decay, Phys. Rev. Lett. **103,** 241802 (2009).

[17] S. N. Gninenko, A resolution of puzzles from the LSND, KARMEN, and MiniBooNE experiments, Phys. Rev. D **83,** 015015 (2011).

[18] S. N. Gninenko, New limits on radiative sterile neutrino decays from a search for single photons in neutrino interactions, Phys. Lett. B **710,** 86 (2012).

[19] M. Masip, P. Masjuan, and D. Meloni, Heavy neutrino decays at MiniBooNE, J. High Energy Phys. 01 (2013) 106.

[20] A. Radionov, Constraints on electromagnetic properties of sterile neutrinos from MiniBooNE results, Phys. Rev. D **88,** 015016 (2013).

[21] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, Dipole portal to heavy neutral leptons, Phys. Rev. D **98,** 115015 (2018).

[22] S. Vergani, N. W. Kamp, A. Diaz, C. A. Argüelles, J. M. Conrad, M. H. Shaevitz, and M. A. Uchida, Explaining the MiniBooNE excess through a mixed model of neutrino oscillation and decay, Phys. Rev. D **104,** 095005 (2021).

[23] L. Alvarez-Ruso and E. Saul-Sala, Neutrino interactions with matter and the MiniBooNE anomaly, Eur. Phys. J. Spec. Top. **230,** 4373 (2021).

[24] P. B. Denton, Y. Farzan, and I. M. Shoemaker, Activating the fourth neutrino of the $3+1$ scheme, Phys. Rev. D **99,** 035003 (2019).

[25] E. Bertuzzo, S. Jana, P. A. N. Machado, and R. Zukanovich Funchal, Dark Neutrino Portal to Explain MiniBooNE excess, Phys. Rev. Lett. **121,** 241801 (2018).

[26] E. Bertuzzo, S. Jana, P. A. N. Machado, and R. Zukanovich Funchal, Neutrino masses and mixings dynamically generated by a light dark sector, Phys. Lett. B **791,** 210 (2019).

[27] P. Ballett, S. Pascoli, and M. Ross-Lonergan, U(1)' mediated decays of heavy sterile neutrinos in MiniBooNE, Phys. Rev. D **99,** 071701 (2019).

[28] C. A. Argüelles, M. Hostert, and Y.-D. Tsai, Testing New Physics Explanations of the MiniBooNE Anomaly at Neutrino Scattering Experiments, Phys. Rev. Lett. **123,** 261801 (2019).

[29] P. Ballett, M. Hostert, and S. Pascoli, Dark neutrinos and a three portal connection to the Standard Model, Phys. Rev. D **101,** 115025 (2020).

[30] P. Ballett, M. Hostert, and S. Pascoli, Neutrino masses from a dark neutrino sector below the electroweak scale, Phys. Rev. D **99,** 091701 (2019).

[31] A. Abdullahi, M. Hostert, and S. Pascoli, A dark seesaw solution to low energy anomalies: MiniBooNE, the muon $(g-2)$, and *BABAR*, Phys. Lett. B **820,** 136531 (2021).

[32] A. Datta, S. Kamali, and D. Marfatia, Dark sector origin of the KOTO and MiniBooNE anomalies, Phys. Lett. B **807,** 135579 (2020).

[33] B. Dutta, S. Ghosh, and T. Li, Explaining $(g-2)_{\mu,e}$, the KOTO anomaly and the MiniBooNE excess in an extended Higgs model with sterile neutrinos, Phys. Rev. D **102,** 055017 (2020).

[34] W. Abdallah, R. Gandhi, and S. Roy, Understanding the MiniBooNE and the muon and electron $g-2$ anomalies with a light Z' and a second Higgs doublet, J. High Energy Phys. 12 (2020) 188.

[35] W. Abdallah, R. Gandhi, and S. Roy, Two-Higgs doublet solution to the LSND, MiniBooNE and muon g-2 anomalies, Phys. Rev. D **104,** 055028 (2021).

[36] A. Hammad, A. Rashed, and S. Moretti, The dark Z' and sterile neutrinos behind current anomalies, Phys. Lett. B **827,** 136945 (2022).

[37] B. Dutta, D. Kim, A. Thompson, R. T. Thornton, and R. G. Van de Water, Solutions to the MiniBooNE Anomaly from New Physics in Charged Meson Decays, Phys. Rev. Lett. **129,** 111803 (2022).

[38] O. Fischer, A. Hernández-Cabezudo, and T. Schwetz, Explaining the MiniBooNE excess by a decaying sterile neutrino with mass in the 250 MeV range, Phys. Rev. D **101,** 075045 (2020).

[39] C.-H. V. Chang, C.-R. Chen, S.-Y. Ho, and S.-Y. Tseng, Explaining the MiniBooNE anomalous excess via a leptophilic ALP-sterile neutrino coupling, Phys. Rev. D **104,** 015030 (2021).

[40] W. Abdallah, R. Gandhi, and S. Roy, Requirements on common solutions to the LSND and MiniBooNE excesses: A post-MicroBooNE study, J. High Energy Phys. 06 (2022) 160.

[41] K. Abe *et al.* (T2K Collaboration), Search for heavy neutrinos with the T2K near detector ND280, Phys. Rev. D **100**, 052006 (2019).

[42] W. Buchmuller, C. Greub, and P. Minkowski, Neutrino masses, neutral vector bosons and the scale of B-L breaking, Phys. Lett. B **267**, 395 (1991).

[43] S. Khalil, Low scale $B - L$ extension of the standard model at the LHC, J. Phys. G **35**, 055001 (2008).

[44] P. Fileviez Perez, T. Han, and T. Li, Testability of type I seesaw at the CERN LHC: Revealing the existence of the $B - L$ symmetry, Phys. Rev. D **80**, 073015 (2009).

[45] S. Khalil, TeV-scale gauged $B - L$ symmetry with inverse seesaw mechanism, Phys. Rev. D **82**, 077702 (2010).

[46] R. Harnik, J. Kopp, and P. A. N. Machado, Exploring ν signals in dark matter detectors, J. Cosmol. Astropart. Phys. 07 (2012) 026.

[47] B. Batell, M. Pospelov, and B. Shuve, Shedding light on neutrino masses with dark forces, J. High Energy Phys. 08 (2016) 052.

[48] C. O. Dib, G. R. Moreno, and N. A. Neill, Neutrinos with a linear seesaw mechanism in a scenario of gauged B-L symmetry, Phys. Rev. D **90**, 113003 (2014).

[49] V. De Romeri, E. Fernandez-Martinez, J. Gehrlein, P. A. N. Machado, and V. Niro, Dark matter and the elusive Z' in a dynamical Inverse Seesaw scenario, J. High Energy Phys. 10 (2017) 169.

[50] D. A. Camargo, M. D. Campos, T. B. de Melo, and F. S. Queiroz, A two Higgs doublet model for dark matter and neutrino masses, Phys. Lett. B **795**, 319 (2019).

[51] C. Argüelles, P. Coloma, P. Hernández, and V. Muñoz, Searches for atmospheric long-lived particles, J. High Energy Phys. 02 (2020) 190.

[52] M. Pospelov, Neutrino physics with dark matter experiments and the signature of new baryonic neutral currents, Phys. Rev. D **84**, 085008 (2011).

[53] S. Baek, H. Okada, and K. Yagyu, Flavour dependent gauged radiative neutrino mass model, J. High Energy Phys. 04 (2015) 049.

[54] A. Kamada and H.-B. Yu, Coherent propagation of PeV neutrinos and the dip in the neutrino spectrum at IceCube, Phys. Rev. D **92**, 113004 (2015).

[55] S. Baek, P. Ko, and W.-I. Park, Singlet portal extensions of the standard seesaw models to a dark sector with local dark symmetry, J. High Energy Phys. 07 (2013) 013.

[56] H. Okada and K. Yagyu, Radiative generation of lepton masses with the U(1)' gauge symmetry, Phys. Rev. D **90**, 035019 (2014).

[57] P. Ko and Y. Tang, $\nu\Lambda$MDM: A model for sterile neutrino and dark matter reconciles cosmological and neutrino oscillation data after BICEP2, Phys. Lett. B **739**, 62 (2014).

[58] C. E. Diaz, S. F. Mantilla, and R. Martinez, Radiative neutrino masses with forbidden seesaw mechanism in an abelian extension of the standard model, arXiv:1712.05433.

[59] T. Nomura and H. Okada, Linear seesaw model with hidden gauge symmetry, Phys. Rev. D **99**, 055033 (2019).

[60] C. Hagedorn, J. Herrero-García, E. Molinaro, and M. A. Schmidt, Phenomenology of the generalised scotogenic model with fermionic dark matter, J. High Energy Phys. 11 (2018) 103.

[61] B. Shakya and J. D. Wells, Exotic sterile neutrinos and pseudo-Goldstone phenomenology, J. High Energy Phys. 02 (2019) 174.

[62] C. W. De Jager, H. De Vries, and C. De Vries, Nuclear charge and magnetization density distribution parameters from elastic electron scattering, At. Data Nucl. Data Tables **14**, 479 (1974); **16**, 580(E) (1975).

[63] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for an Excess of Electron Neutrino Interactions in MicroBooNE Using Multiple Final-State Topologies, Phys. Rev. Lett. **128**, 241801 (2022).

[64] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for an anomalous excess of inclusive charged-current ν_e interactions in the MicroBooNE experiment using Wire-Cell reconstruction, Phys. Rev. D **105**, 112005 (2022).

[65] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for an anomalous excess of charged-current quasielastic νe interactions with the MicroBooNE experiment using deep-learning-based reconstruction, Phys. Rev. D **105**, 112003 (2022).

[66] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for an anomalous excess of charged-current νe interactions without pions in the final state with the MicroBooNE experiment, Phys. Rev. D **105**, 112004 (2022).

[67] C. A. Argüelles, I. Esteban, M. Hostert, K. J. Kelly, J. Kopp, P. A. N. Machado, I. Martinez-Soler, and Y. F. Perez-Gonzalez, MicroBooNE and the νe Interpretation of the MiniBooNE Low-Energy Excess, Phys. Rev. Lett. **128**, 241802 (2022).

[68] P. B. Denton, Sterile Neutrino Search with MicroBooNE's Electron Neutrino Disappearance Data, Phys. Rev. Lett. **129**, 061801 (2022).

[69] V. Brdar, O. Fischer, and A. Y. Smirnov, Model-independent bounds on the nonoscillatory explanations of the MiniBooNE excess, Phys. Rev. D **103**, 075008 (2021).

[70] K. Abe *et al.* (T2K Collaboration), The T2K experiment, Nucl. Instrum. Methods Phys. Res., Sect. A **659**, 106 (2011).

[71] K. Abe *et al.* (T2K Collaboration), Measurement of the charged-current electron (anti-)neutrino inclusive cross-sections at the T2K off-axis near detector ND280, J. High Energy Phys. 10 (2020) 114.

[72] S. Assylbekov *et al.*, The T2K ND280 off-axis pi-zero detector, Nucl. Instrum. Methods Phys. Res., Sect. A **686**, 48 (2012).

[73] K. Gilje, Neutral current π^0 production rate measurement on-water using the π^0 detector in the near detector of the T2K experiment, Ph.D. thesis, SUNY, Stony Brook, 2014.

[74] M. Lamoureux, Recherche de neutrinos lourds avec l'expérience T2K, Ph.D. thesis, Saclay, 2018.

[75] C. A. Argüelles, N. Foppiani, and M. Hostert, Heavy neutral leptons below the kaon mass at hodoscopic neutrino detectors, Phys. Rev. D **105**, 095006 (2022).

[76] A. M. Abdullahi, J. Hoefken Zink, M. Hostert, D. Massaro, and S. Pascoli, DarkNews: A Python-based event generator for heavy neutral lepton production in neutrino-nucleus scattering, arXiv:2207.04137.

[77] K. Abe *et al.* (T2K Collaboration), T2K ND280 upgrade—technical design report, arXiv:1901.03750.

[78] K. Abe *et al.*, Proposal for an extended run of T2K to 20 × 10^{21} POT, arXiv:1609.04111.

[79] S. King, G. Christodoulou, M. Lamoureux, and T. T. Experiment, Data release for the "Measurement of the charged-current electron (anti-)neutrino inclusive cross-sections at the T2K off-axis near detector ND280", Zenodo, 10.5281/zenodo.5543856 (2021).

[80] A. Buckley, H. Hoeth, H. Lacker, H. Schulz, and J. E. von Seggern, Systematic event generator tuning for the LHC, Eur. Phys. J. C **65**, 331 (2010).

[81] M. Krishnamoorthy, H. Schulz, X. Ju, W. Wang, S. Leyffer, Z. Marshall, S. Mrenna, J. Müller, and J. B. Kowalkowski, Apprentice for event generator tuning, EPJ Web Conf. **251**, 03060 (2021).

[82] D. G. Cerdeño, A. Cheek, E. Reid, and H. Schulz, Surrogate models for direct dark matter detection, J. Cosmol. Astropart. Phys. 08 (2018) 011.

[83] M. G. Aartsen *et al.* (IceCube Collaboration), Efficient propagation of systematic uncertainties from calibration to analysis with the SnowStorm method in IceCube, J. Cosmol. Astropart. Phys. 10 (2019) 048.

[84] G. Peter Lepage, A new algorithm for adaptive multidimensional integration, J. Comput. Phys. **27**, 192 (1978).

[85] G. P. Lepage, Adaptive multidimensional integration: VEGAS enhanced, J. Comput. Phys. **439**, 110386 (2021).

[86] P. Lepage, gplepage/vegas: vegas version 5.1.1 (2022).

[87] C. A. Argüelles, A. Schneider, and T. Yuan, A binned likelihood for stochastic models, J. High Energy Phys. 06 (2019) 030.

[88] P. Ilten, Y. Soreq, M. Williams, and W. Xue, Serendipity in dark photon searches, J. High Energy Phys. 06 (2018) 004.

[89] T. Asaka, S. Eijima, and A. Watanabe, Heavy neutrino search in accelerator-based experiments, J. High Energy Phys. 03 (2013) 125.

[90] A. de Gouvêa, P. J. Fox, R. Harnik, K. J. Kelly, and Y. Zhang, Dark tridents at off-axis liquid argon neutrino detectors, J. High Energy Phys. 01 (2019) 001.

[91] J. Isaacson, W. I. Jay, A. Lovato, P. A. N. Machado, and N. Rocco, New approach to intranuclear cascades with quantum Monte Carlo configurations, Phys. Rev. C **103**, 015502 (2021).

[92] J. Isaacson, S. Höche, D. Gutierrez Lopez, and N. Rocco, Novel event generator for the automated simulation of neutrino scattering, Phys. Rev. D **105**, 096006 (2022).

[93] J. M. Campbell *et al.*, Event generators for high-energy physics experiments, arXiv:2203.11110.

[94] M. Antonello *et al.* (MicroBooNE, LAr1-ND, and ICARUS-WA104 Collaborations), A proposal for a three detector short-baseline neutrino oscillation program in the fermilab booster neutrino beam, arXiv:1503.01520.

[95] P. A. Machado, O. Palamara, and D. W. Schmitz, The short-baseline neutrino program at Fermilab, Annu. Rev. Nucl. Part. Sci. **69**, 363 (2019).

[96] S. Adrian-Martinez *et al.* (KM3Net Collaboration), Letter of intent for KM3NeT 2.0, J. Phys. G **43**, 084001 (2016).

[97] B. Abi *et al.* (DUNE Collaboration), Deep underground neutrino experiment (DUNE), far detector technical design report, volume II: DUNE physics, arXiv:2002.03005.

[98] K. Abe *et al.* (Hyper-Kamiokande Collaboration), Hyper-kamiokande design report, arXiv:1805.04163.