

Searching for physics beyond the Standard Model in an off-axis DUNE near detector

Moritz Breitbach,^a Luca Buonocore,^b Claudia Frugiuele,^c Joachim Kopp^{d,a}
and Lukas Mittnacht^a

^a*PRISMA Cluster of Excellence and Mainz Institute for Theoretical Physics,
Johannes Gutenberg-Universität,
Staudingerweg 7, Mainz, Germany*

^b*Physik Institut, Universität Zürich,
Zürich, Switzerland*

^c*INFN, Sezione di Milano,
Via Celoria 16, I-20133, Milano, Italy*

^d*Theoretical Physics Department, CERN,
Geneva, Switzerland*

E-mail: breitbach@uni-mainz.de, lbuono@physik.uzh.ch,
claudia.frugiuele@cern.ch, jkopp@cern.ch, lmittna@uni-mainz.de

ABSTRACT: Next generation neutrino oscillation experiments like DUNE and T2HK are multi-purpose observatories, with a rich physics program beyond oscillation measurements. A special role is played by their near detector facilities, which are particularly well-suited to search for weakly coupled dark sector particles produced in the primary target. In this paper, we demonstrate this by estimating the sensitivity of the DUNE near detectors to the scattering of sub-GeV DM particles and to the decay of sub-GeV sterile neutrinos (“heavy neutral leptons”). We discuss in particular the importance of the DUNE-PRISM design, which allows some of the near detectors to be moved away from the beam axis. At such off-axis locations, the signal-to-background ratio improves for many new physics searches. We find that this leads to a dramatic boost in the sensitivity to boosted DM particles interacting mainly with hadrons, while for boosted DM interacting with leptons, data taken on-axis leads to marginally stronger exclusion limits. Searches for heavy neutral leptons perform equally well in both configurations.

KEYWORDS: Beyond Standard Model, Neutrino Physics

ARXIV EPRINT: [2102.03383](https://arxiv.org/abs/2102.03383)

Contents

1	Introduction	1
2	Light dark matter interacting via a dark photon	3
2.1	Dark matter production and detection	4
2.2	Backgrounds	7
2.3	Statistical analysis	9
2.4	Sensitivity to light dark matter in the dark photon model	9
3	Leptophobic dark matter	10
3.1	Dark matter production and detection	11
3.2	Backgrounds	13
3.3	Sensitivity to leptophobic dark matter	14
4	Heavy neutral leptons	15
4.1	Production of heavy neutral leptons	16
4.2	Heavy neutral lepton decay	17
4.3	Backgrounds and analysis cuts	18
4.4	Sensitivity to heavy neutral leptons	21
5	Summary and conclusions	25
A	Simulation of dark matter production in the dark photon model	26
B	Statistical analysis	28

1 Introduction

The near detectors of long-baseline neutrino experiments, once considered an afterthought to reduce systematic uncertainties in oscillation measurements, are nowadays independent experiments in their own right. Besides delivering a wealth of data on neutrino interaction physics, it has been realized that they could also serve to probe the existence of physics beyond the Standard Model (SM) [1] such as heavy neutral leptons or other new particles, possibly connected to the dark matter (DM) puzzle [2–16]. Indeed, signatures of new dark particles at these experiments can be linked in a predictive way to compelling scenarios of light DM. For instance, “invisible” decays of new sub-GeV particles that mediate light DM-SM interactions, can be searched for at neutrino fixed target facilities by looking for scattering of the decay products off nucleons and/or electrons in the near detector. A light DM program at neutrino facilities could complement the next generation light DM direct detection program [17], in particular, the upcoming experiment SENSEI [18]. So far, only

the Fermilab-based neutrino experiment MiniBooNE has performed dedicated searches for light DM [19, 20] and very recently MicroBooNE has released the first search for heavy neutral leptons (HNL) [21] and for a light Higgs decaying into e^+e^- [22], but the untapped potential is big, with many past and ongoing experiments having the capability to supersede MiniBooNE’s sensitivity [7, 8, 10, 14]. Importantly, these searches can typically be done fully parasitically to the main neutrino program [8, 10].

The near detector physics program will be taken to the next level by the DUNE-PRISM detectors, to be installed 574 m downstream from the target [23] at the long-baseline neutrino facility (LBNF) at Fermilab, the neutrino source for the DUNE experiment. These detectors — a liquid argon time projection chamber (TPC) and a magnetized gaseous argon TPC — will be mounted on a movable platform, allowing them to be displaced up to 30.5 m (53 mrad) away from the beam axis. This capability mainly serves the detectors’ primary purpose, namely constraining the unoscillated neutrino flux and measuring the neutrino cross sections. In particular, the neutrino spectrum changes as a function of the off-axis angle (for kinematic reasons), while the cross sections obviously do not. Therefore, taking data at different off-axis positions will allow DUNE-PRISM to disentangle the uncertainties in the neutrino spectrum from the uncertainties in the neutrino cross section.

In this paper, we will discuss the impact of the DUNE-PRISM concept on searches for physics beyond the SM. More specifically, we will consider the production of light (\lesssim GeV) and very weakly interacting new particles in the target, followed by their interaction or decay inside the DUNE-PRISM detectors.

Among the numerous extensions of the SM that can be probed in DUNE-PRISM and other accelerator neutrino experiments, we will consider in particular: (1) light (\lesssim GeV) DM particles produced with a large Lorentz boost and detectable via dark photon-mediated DM-electron scattering; (2) light DM particles detectable via DM-nucleus scattering (leptophobic DM); and (3) heavy neutral leptons (sterile neutrinos) decaying to various combinations of neutrinos, charged leptons, and hadrons. In the following sections, we will introduce these scenarios one by one and discuss the anticipated sensitivity of DUNE-PRISM, both on-axis and off-axis. Specifically, section 2 will be focused on dark photon-mediated DM, section 3 will deal with leptophobic DM, and section 4 will be about heavy neutral leptons. We will discuss our findings and conclude in section 5. Let us comment that similar searches for DUNE have been considered previously in refs. [8, 16, 24, 25] for the case of on-axis detectors, and for additional data taking away from the beam axis in ref. [13].

We will go beyond these studies in two important ways:

- We will investigate the usefulness of taking DUNE off-axis data for additional scenarios, for which this has never been done before. In particular, we will consider scattering of light, leptophobic DM and decays of heavy neutral leptons.
- We will also reconsider DM scattering on electrons, previously studied in ref. [13]. As a cross-check, we will reproduce the total rates analysis carried out in this reference, but we will also show that an analysis including the electron recoil spectrum is equally sensitive on-axis and off-axis.

2 Light dark matter interacting via a dark photon

Most direct searches for DM lose sensitivity at DM masses below a few GeV, motivating a new experimental program that focuses specifically on this mass range [17].

One of the simplest and most generic models for DM in the MeV-GeV mass range augments the SM by a scalar DM particle ϕ (or a Majorana fermion) and a new U(1)' gauge boson, A' . The relevant terms in the Lagrangian read

$$\mathcal{L}_{\text{DM}} = \mathcal{L}_{A'} + \mathcal{L}_{\phi}, \quad (2.1)$$

with

$$\mathcal{L}_{A'} = -\frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{m_{A'}^2}{2}A'^{\mu}A'_{\mu} - \frac{1}{2}\epsilon F'_{\mu\nu}F^{\mu\nu}, \quad (2.2)$$

and

$$\mathcal{L}_{\phi} = ig' A'^{\mu} J_{\mu}^{\phi} + (\partial_{\mu}\phi^{\dagger})(\partial^{\mu}\phi) - m_{\phi}^2\phi^{\dagger}\phi, \quad (2.3)$$

where $J_{\mu}^{\phi} = [(\partial_{\mu}\phi^{\dagger})\phi - \phi^{\dagger}(\partial_{\mu}\phi)]$ is the DM current, g' is the U(1)' gauge coupling, $F_{\mu\nu}$ and $F'_{\mu\nu}$ are the U(1) and U(1)' field strength tensors, respectively, and ϵ parameterizes the small kinetic mixing between the dark and visible photons. ϵ thus ultimately controls the interaction strength between the dark photon and SM particles. For DM lighter than half the dark photon mass ($m_{\phi} < m_{A'}/2$), the thermal relic abundance of ϕ is determined by its annihilation cross section to SM fermions,

$$\sigma(\phi\phi^{\dagger} \rightarrow f\bar{f})v_{\text{rel}} \sim \frac{8\pi v_{\text{rel}}^2 Y}{m_{\phi}^2}, \quad (2.4)$$

where we have defined the effective coupling strength

$$Y \equiv \epsilon^2 \alpha_D \left(\frac{m_{\phi}}{m_{A'}}\right)^4. \quad (2.5)$$

As usual, v_{rel} denotes the relative velocity of the two annihilating DM particles, and $\alpha_D \equiv g'^2/(4\pi)$. In the following, we will present our results in the m_{ϕ} - Y plane since this choice makes it easiest to highlight those regions of parameter space where the correct DM thermal abundance is obtained [17, 26].

One important feature of the annihilation cross section in eq. (2.4) is its v_{rel}^2 suppression, through which strong constraints based on precise measurements of the temperature anisotropies of the cosmic microwave background radiation [27–30] are avoided. For models with unsuppressed annihilation, these constraints would rule out thermal freeze-out production of DM candidates with a mass below ~ 10 GeV. Thanks to the velocity suppression, scalars or Majorana fermions can account for the totality of the DM abundance via thermal freeze-out even at masses below 1 GeV.

The kinetic mixing term in eq. (2.2) implies that any process that can create a photon can also create a dark photon, provided this is kinematically allowed. In a meson production

target like the ones employed in neutrino beam experiments, dark photons can be copiously produced in meson decays such as $\pi^0 \rightarrow \gamma A'$ or $\eta \rightarrow \gamma A'$, with a smaller contribution from bremsstrahlung. The dark photon couples to the dark current with coupling strength g' , and to the SM electromagnetic current with coupling strength ϵe . We will consider in particular the case where the dark photon mass, $m_{A'}$, is larger than twice the DM mass, m_ϕ . In this case, any A' produced in the target will rapidly decay, almost exclusively to $\phi\phi^\dagger$. Hence, a beam of ϕ particles will travel alongside the neutrino beam and eventually reach the near detector, where ϕ particles can scatter on nuclei and electrons. It is in particular the latter channel — ϕ -electron scattering — that we will focus on because in this channel neutrino-induced backgrounds are smaller [2, 10, 14].

2.1 Dark matter production and detection

In a proton beam dump, dark photons with masses below ~ 1 GeV are mainly produced in the decays of the lightest neutral mesons, π^0 and η , and in proton bremsstrahlung via the process $pp \rightarrow ppA'$. Production processes induced by leptonic secondary particles and their bremsstrahlung are usually subdominant, even if not completely negligible as reported in a detailed calculation in ref. [31]. We do not consider the latter type of processes in this work, and therefore our estimates should be considered conservative in this regard. In the mass window considered, production mechanisms that can be described in perturbative quantum chromodynamics (QCD), such as Drell-Yan production, are negligible and are not taken into account here. In the simulation of the dark photon signal, we realistically take into account correlations between the geometric acceptance of the detector and the angular spread of the DM flux. For the detector geometry, we consider for simplicity a cylindrical shape oriented along the beam axis with transverse surface given by a circle of radius 3.5 m. While this simplified detector model does not exactly match the envisioned geometry of the DUNE-PRISM detectors, the error introduced by our approximation should be negligible compared to the intrinsic uncertainties of the flux prediction.

Meson decay. The production of dark photons in meson decays occurs mostly via transitions of the form

$$X \rightarrow \gamma A' \rightarrow \gamma \phi \phi^\dagger, \tag{2.6}$$

where $X = \pi^0, \eta$. Reactions involving higher mass mesons are possible but usually subdominant, as shown for example in refs. [32, 33]. The ρ and ω resonances, which lead to a sizeable enhancement of the production rate in a narrow mass region, are effectively taken into account within our formalism for proton bremsstrahlung, as detailed below. The transition in eq. (2.6) can proceed either via an on-shell or an off-shell A' . We assume that the decay is dominated by the on-shell mode and make use of the formula [2, 34]

$$\frac{\text{BR}(X \rightarrow \gamma A')}{\text{BR}(X \rightarrow \gamma \gamma)} \simeq 2 \epsilon^2 \left(1 - \frac{m_{A'}^2}{m_X^2}\right)^3, \quad X = \pi^0, \eta \tag{2.7}$$

for the corresponding branching ratio. This expression has been derived in the narrow width approximation (see ref. [35] for a discussion on the full treatment of off-shell and its implication on the sensitivity).

As a first step towards computing the rate of DM-electron scattering events expected in the DUNE near detectors, we need to model the spectra of π^0 s and η s produced in DUNE’s primary target. This is crucial as any systematic bias in these spectra will propagate through the simulation chain and affect the yield of signal events. This is even more relevant if one considers, as done in this work, the possibility offered by the DUNE-PRISM concept to have a movable detector. We consider two samples of mesons:

- *primary-only (SoftQCD)*, which includes only the mesons produced in the primary interaction of the proton impinging on the target, assuming that all protons eventually interact there. We employ PYTHIA v8.230 [36] and adopt the flag `SoftQCD` as the main mode for the simulation. For the simulation of relatively low-energy fixed target experiments with PYTHIA, the flag `SoftQCD:All` should be preferably used, as reported also in ref. [37].
- *beam-dump*, provided as ancillary material of ref. [31], which includes the interaction of secondary particles propagating in the DUNE target using GEANT4 [38].

The *beam-dump* sample is more realistic and is taken as our default choice to study the DUNE sensitivity. We consider the other sample for sanity checks and for comparison with the result of ref. [13], where PYTHIA was employed for the simulation of mesons as well. We found discrepancies with the latter that we traced back to a different setup in the PYTHIA generator.¹ Further details and technical aspects on the comparison are reported in appendix A.

With the meson fluxes in hand, we use MADDUMP [39] to generate samples of signal events of DM-electron scattering in the DUNE near detectors for all relevant parameter points. The program, once instrumented with a suitable UFO model file which encodes the coupling structure of the DM and the dark photon [40], takes care of all the steps of the simulation chain, from the decay of the parent mesons into DM particles to the detection process in the DUNE near detectors, including the computation of geometric acceptances and other effects due to the finite size of the detector.

In figure 1 we compare the results of the two different signal predictions by plotting the rate of signal + background events in each case as a function of the detector location Δx_{OA} relative to the beam axis. The broad features visible in the figure are consistent with expectations: both the signal and background rates decrease as the detector is moved away from the focus direction of the beam. The signal-to-background ratio, however, generally increases with Δx_{OA} because the nearly massless neutrinos that are responsible for the background inherit more of the forward boost of their parent mesons than the much heavier A' bosons from which the DM originates. Comparing our *primary-only (SoftQCD)* (red dot-dashed) sample to the *beam-dump* one (blue dashed), we observe that they are of similar size in the on-axis bins, but diverge with increasing off-axis angle: the event rate predicted by the *beam-dump* sample drops less rapidly as the detector is moved away from the beam axis. This is indeed consistent with our expectation as there is a strong correlation between

¹We thank the authors of ref. [13] for helping us understand the PYTHIA flags that were used in their simulation.




Figure 1. Expected number of DM-electron scattering events in the DUNE near detectors as a function of the detector position relative to the beam axis for the neutrino mode run. The background (solid gray histogram) corresponds to elastic neutrino and anti-neutrino scattering on electrons. Following ref. [13] we assume that the background from charged current quasi-elastic interactions can be made negligible by applying an energy-dependent cut on the lepton angle, which barely affects the signal and the elastic neutrino-electron scattering events. For the signal + background histograms, we show separately the two signal samples introduced in the main text: the PYTHIA-based sample of primary mesons, *primary-only (SoftQCD)* (red dot-dashed), and the sample *beam-dump* (blue dashed) based on the ancillary material of ref. [31]. See text for a discussion of the differences between these two samples.

the location of the detector and the energy spectrum of the DM particles, which in turn originates from the spectra of the parent mesons. In particular, when the detector is on-axis, the flux it receives is dominated by the decay products of relatively energetic mesons, which are most likely produced in the primary proton-proton interaction. Therefore, this flux is reliably modeled by the *primary-only (SoftQCD)* method. Off-axis, however, the DM energy spectrum is dominated by much lower-energy particles, making it much more sensitive to the decays of soft mesons produced in secondary interactions in the target. The *beam-dump* sample includes these mesons, explaining why the discrepancy between the blue and red curves in figure 1 increases with Δx_{OA} . We conclude that the modeling of secondary interactions is crucial for the off-axis strategy.

Proton bremsstrahlung. In the dark photon mass range $500 \text{ MeV} \lesssim m_{A'} \lesssim 1 \text{ GeV}$, i.e. above the η threshold, proton bremsstrahlung dominates dark photon production. Bremsstrahlung is preferentially emitted in the forward direction (collinear with the incoming proton) and in this limit can be well described by a generalization of the Fermi-Williams-Weizsäcker method [41–43] (or “equivalent photon method”). This method is based on the

assumption that proton-nucleon scattering is dominated by exchange of vector bosons that are close to on-shell. Again we rely on MADDUMP for our simulations, which implements bremsstrahlung following refs. [44, 45]. Let us parameterize the 4-momentum vector of the emitted A' as $p_{A'} = (E_{A'}, p_T \cos(\phi), p_T \sin(\phi), zP)$, with $E_{A'} \simeq zP + (p_T^2 + m_{A'}^2)/(2zP)$. Here, P is the momentum of the incident proton, z is the fraction of the proton momentum carried by the outgoing A' , p_T is the momentum perpendicular to the beam momentum, and ϕ is the azimuthal angle. We generate unweighted A' events according to the differential production rate

$$\frac{d^2 N_{A'}}{dz dp_T^2} = \frac{\sigma_{pA}(s')}{\sigma_{pA}(s)} F_{1,p}^2(m_{A'}^2) w_{ba}(z, p_T^2), \quad (2.8)$$

where $\sigma_{pA}(s)$ denotes the total interaction cross section of the incoming protons with a target nucleus of mass number A , $s = 2m_p E_p$ is the square of the center-of-mass energy, and $s' = 2m_p(E_p - E_{A'})$. The ratio of cross sections $\sigma_{pA}(s')/\sigma_{pA}(s)$ compares the probability that the incoming proton interacts after having emitted a photon to its total interaction probability. For the proton form factor $F_{1,p}(m_{A'}^2)$, we use the parameterization from ref. [46] in the time-like region, so that off-shell mixing with vector mesons such as the ρ and ω is effectively included in our calculation, leading to a resonance peak in the A' production rate around $m_{A'} \simeq 770$ MeV [47]. Finally, the photon splitting function is

$$w_{ba}(z, p_T^2) = \frac{\epsilon^2 \alpha}{2\pi H} \left[\frac{1 + (1-z)^2}{z} - 2z(1-z) \left(\frac{2m_p^2 + m_{A'}^2}{H} - z^2 \frac{2m_p^4}{H^2} \right) + 2z(1-z)[1 + (1-z)^2] \frac{m_p^2 m_{A'}^2}{H^2} + 2z(1-z)^2 \frac{m_{A'}^4}{H^2} \right], \quad (2.9)$$

with $H = p_T^2 + (1-z)m_{A'}^2 + z^2 m_p^2$. The number of produced A' events is estimated by integrating eq. (2.8) over a region well within the realm of the collinear approximation, i.e. a region where the kinematic conditions

$$E_p, E_{A'}, E_p - E_{A'} \gg m_p, m_{A'}, p_T \quad (2.10)$$

hold. In particular, following refs. [44, 45, 48], we use the integration intervals $z \in [0.1, 0.9]$ and $p_T < 1$ GeV.

2.2 Backgrounds

As the experimental signature of $\phi-e^-$ scattering is a single, energetic recoil electron, the main backgrounds to this search in a neutrino beam experiment are due to neutrino-electron scattering and, if the final state hadronic system stays unidentified, charged current (CC) ν_e -nucleon interactions. We estimate the backgrounds using GENIE v3.00.06 [49]. In particular, we simulate the relevant processes — $\nu_e e^-$ scattering and CC ν_e interactions on nuclei — in argon. The resulting event rates are then weighted bin-by-bin by the various on- and off-axis fluxes. The simulated DUNE neutrino fluxes have been taken from ref. [50], and extrapolated to higher energies, where Monte Carlo statistics in the simulations from

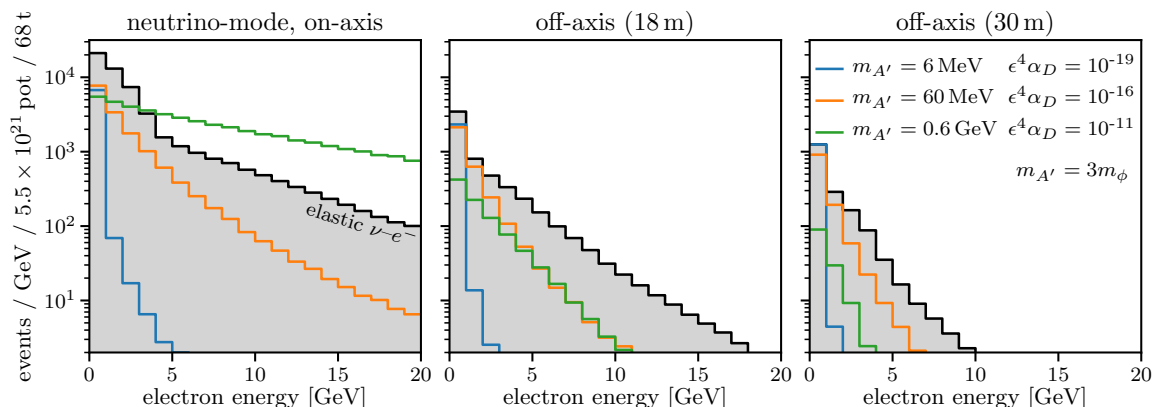


Figure 2. Signal and background spectra in DUNE-PRISM for the dark photon model after 5 years of data taking (5.5×10^{21} pot) in neutrino mode. Colored histograms show the DM-electron scattering signal for different exemplary sets of model parameter points, imposing $m_{A'} = 3m_\phi$. The gray shaded background histogram is based on a simulation of elastic neutrino-electron scattering. Comparing data taken all on-axis (left panel) to data taken at 18 m (31.36 mrad) off-axis (center panel) or 30 m (52.26 mrad) off-axis (right panel), we see that, as expected, both the signal and background rates are significantly lower in the off-axis positions, especially at high energy. At low energies, where most of the events are concentrated, the signal-to-background ratio becomes better when going off-axis.

ref. [50] is too low. We perform this extrapolation by fitting the tail of the available fluxes linearly in log-space. The expected flux at very large neutrino energies, where the existing fluxes lack statistics, is then obtained by evaluating the fit function.

From the kinematics of elastic scattering, we know that both the signal process $\phi e^- \rightarrow \phi e^-$ and the background process $\nu e^- \rightarrow \nu e^-$ obey $E_e \theta_e^2 < 2m_e$, with E_e (θ_e) being the final state electron's total energy (scattering angle). Neutrino-nucleon interactions, on the other hand, lead to larger scattering angles. Since the expected angular resolution of DUNE-PRISM is sufficient [51], we impose the above condition as a kinematic cut, which leaves us with a comparatively small number of events involving neutrino-nucleon interactions. We therefore neglect the latter and only consider the elastic neutrino-electron scattering contribution in our analyses.

In figure 2, we plot the projected electron energy spectra for both the DM-electron scattering signal in the dark photon model and for the neutrino-electron scattering background. We see that both the signal and the background peak at the lowest energies, but feature a long tail towards higher energies. For the signal, the tail is most pronounced for heavier A' , corresponding to heavier DM particles, as we impose $m_{A'} = 3m_\phi$ here. Heavy A' are produced when heavy mesons decay in the LBNF target and are therefore only kinematically accessible in very hard proton-proton collisions. The high-energy tails of both the signal and the background are strongly suppressed off-axis because the production of very energetic DM particles (for the signal) and neutrinos (for the background) occurs in production events that are strongly boosted in the forward direction.

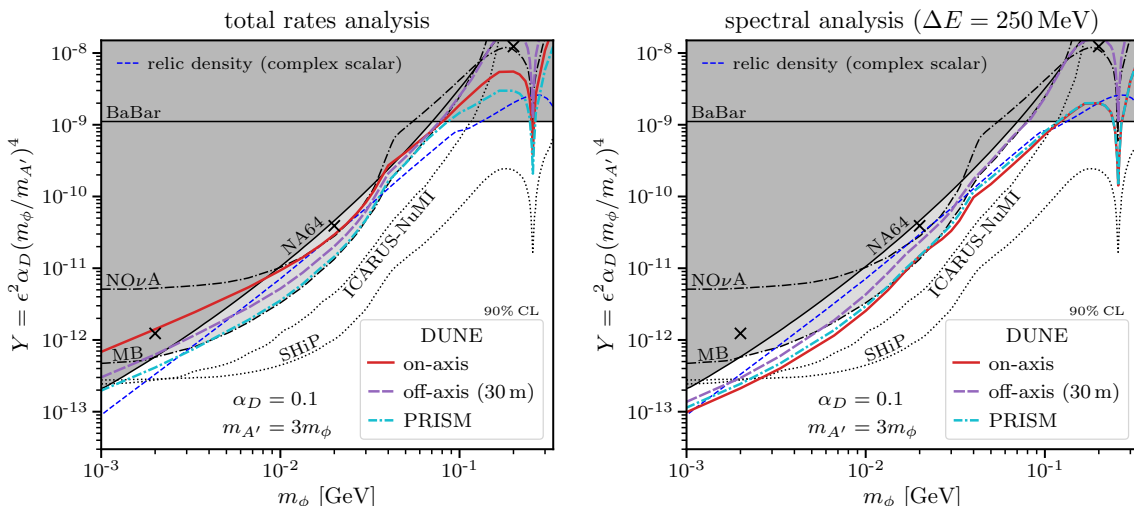


Figure 3. Expected upper exclusion limits for the dark photon model, assuming a total DUNE-PRISM running time of 5 years (5.5×10^{21} pot) in neutrino mode. We compare results for on-axis-only running (red solid), off-axis-only running (purple dashed), and a realistic DUNE-PRISM strategy with equal amounts of data taken at seven different locations (cyan dot-dashed). The left panel corresponds to a total rates (cut & count) analysis as in ref. [13], while the analysis in the right panel includes spectral information. Gray shaded regions indicate existing limits on the model, while dotted black lines show projections for other future experiments. Black crosses indicate the exemplary model parameter points presented in figure 2.

2.3 Statistical analysis

To derive sensitivity limits from our predicted signal and background rates, we use standard frequentist techniques. In particular, we test the signal + background hypothesis against simulated background-only data. To do so, we use a Poissonian log-likelihood function, $\log \mathcal{L}(\Theta, \mathbf{X})$, defined in eq. (B.1), which depends on the physical model parameters $\Theta = (\epsilon^4 \alpha_D, m_{A'}, m_\phi)$ and a set of nuisance parameters \mathbf{X} . The latter parameterize systematic normalization uncertainties and spectral “tilts” in both the signal and the background spectra. We consider systematic errors that are uncorrelated between different on-/off-axis positions (assuming 1% relative error) in addition to errors which are correlated among all positions (assuming 10% relative error). The sensitivity limits on $\epsilon^4 \alpha_D$ for fixed values of $m_{A'}$ and m_ϕ are determined by comparing the log-likelihood ratio, defined in eq. (B.4), to the 90% quantile of a χ^2 distribution with one degree of freedom. See appendix B for further details on our statistical procedure.

2.4 Sensitivity to light dark matter in the dark photon model

We present our main results for the dark photon model in figure 3. In this figure, we compare on the one hand different running strategies: all data taken on-axis, all data taken off-axis, and combining data taken at different on-axis and off-axis locations as in DUNE-PRISM. On the other hand, we also compare two different analysis strategies, namely a total rates analysis ($n_{\text{bins}} = 1$ in eq. (B.1)) in the left panel and a spectral

analysis ($n_{\text{bins}} = 80$ equal-width bins up to 20 GeV) in the right panel. The former type of analysis is similar to the one discussed in ref. [13], and we confirm the main conclusion of these authors, namely that the DUNE-PRISM strategy of combining runs in seven different on-axis and off-axis locations (cyan dot-dashed) benefits the sensitivity to light DM in the dark photon model. It yields better results than both an on-axis-only run (red solid) and an off-axis-only run (purple dashed). Interestingly, though, we reach a different conclusion when including the event spectrum: as shown in the right panel of figure 3, the sensitivity in this case is about the same for on-axis-only running and for the PRISM strategy. This can be understood by going back to figure 2, where we see that the signal-to-background ratio at energies $\gtrsim 2$ GeV is significantly better on-axis than it is off-axis. For a total rates analysis, these high energy events do not contribute because of the steep drop of the event rate compared to the lowest energy bins. A spectral analysis, however, is able to harness also the statistical power of high-energy events and therefore suffers more from the poorer signal-to-background ratio in the off-axis position.

The overall shape of the exclusion curves in figure 3 can be understood mostly from the ϕ production rate, which drops at larger masses, where fewer production modes are available. The spectral feature at $m_\phi \sim 230$ MeV, corresponding to $m_{A'} \sim 700$ MeV is related to the ρ resonance: when $m_{A'} = m_\rho$, dark photons and ρ vector mesons exhibit maximal mixing, leading to very efficient A' production and thus strong limits.

We compare the DUNE sensitivity to existing limits from BaBar [52] and NA64 [53], to a recast of NuMI off-axis data from NO ν A [10] and MiniBooNE (MB) [14], and to the expected sensitivities of ICARUS-NuMI off-axis [14] and SHiP [33]. Note that we have rescaled the ICARUS-NuMI limit to an integrated luminosity of 2.5×10^{21} protons on target (pot), corresponding to 5 years of NuMI running at the nominal beam power of 700 kW. We choose to present here some recasts that have not been officially approved by the respective experimental collaboration and for this reason have been omitted in some previous studies such as the ‘‘Physics Beyond Colliders’’ study at CERN [54]. We include such unofficial recasts only for neutrino experiments like NO ν A [10], but not for other experiments like E137 [31, 55] or BEBC [14]. Also we do not present projections for future proposed experiments other than SHiP for the same reason, for a summary see ref. [54]. We see that DUNE-PRISM can probe important new regions of parameter space, with a sensitivity to $Y = \epsilon^2 \alpha_D (m_\phi/m_{A'})^4$ that is up to half an order of magnitude better than existing constraints for some m_ϕ . Moreover, DUNE-PRISM can improve on the projected sensitivity for ICARUS-NuMI both in the small mass region ($m_{A'} \lesssim 10$ MeV) and at large dark photon mass ($m_{A'} \gtrsim 100$ MeV).

3 Leptophobic dark matter

The second scenario we are going to consider in this paper is a leptophobic DM model, where the visible and the dark sector interact via a new gauge interaction under which the leptons are neutral. As a concrete example, we will consider the force associated with a gauged baryon number symmetry $U(1)_B$. We will call the leptophobic gauge field Z' and the corresponding gauge coupling g_Z . The relevant terms in the Lagrangian of the

leptophobic model are then

$$\mathcal{L}_{\text{leptophobic}} \supset ig_Z z_\phi Z'^\mu J_\mu^\phi + \partial_\mu \phi^\dagger \partial^\mu \phi - m_\phi^2 \phi^\dagger \phi + g_Z z_q \sum_q \bar{q} \gamma^\mu q Z'_\mu, \quad (3.1)$$

where $J_\mu^\phi = [(\partial_\mu \phi^\dagger)\phi - \phi^\dagger(\partial_\mu \phi)]$. The sum in the last term runs over all quark flavors, and z_ϕ, z_q denote the $U(1)_B$ charges of DM and of SM quarks, respectively. The leptophobic nature of the Z' makes this particle rather elusive to experimental probes, as we will see in the following section. However, strong theoretical constraints arise from the fact that the $U(1)_B$ gauge symmetry is anomalous, so that extra fermions need to be added to the model to cancel anomalies [56–59]. Since these constraints are somewhat dependent on the ultraviolet completion of the model, we will not include them in our sensitivity plots.

3.1 Dark matter production and detection

We focus on Z' candidates with masses $m_{Z'} \geq 2 \text{ GeV}$. In this range, the dominant production mechanism is given by prompt production, which can be described by standard methods in perturbative QCD.

The number of DM particles produced in the collision of primary protons impinging on the carbon target is given by the formula

$$N_\phi = 2N_{\text{POT}} \frac{\sigma_{pA \rightarrow \phi\phi^\dagger}}{\sigma_{pA, \text{tot}}} = A^{0.29} N_{\text{POT}} \frac{\sigma_{pN \rightarrow \phi\phi^\dagger}}{\sigma_{pN, \text{tot}}} \quad (3.2)$$

where N_{POT} is the number of protons on target, $A = 12$ is the mass number of carbon (the target material used in the LBNF), and σ_{pA} and σ_{pN} stand for the proton-nucleus and proton-nucleon cross sections, respectively. In the above, we have assumed linear scaling with A for the cross section of DM pair production, $\sigma_{pA \rightarrow \phi\phi^\dagger} = A\sigma_{pN \rightarrow \phi\phi^\dagger}$, and an effective total cross section per nucleon $\sigma_{pA, \text{tot}} = A^{0.71}\sigma_{pN, \text{tot}}$ for proton-carbon collisions [60]. We adopt the approximation $\sigma_{pN, \text{tot}} = 40 \text{ mb}$ [61] and compute $\sigma_{pA \rightarrow \phi\phi^\dagger}$ numerically in MAD-DUMP, using a UFO implementation of the leptophobic model as done in ref. [39]. The computation is carried out at tree-level in QCD according to the standard parton model formula

$$\sigma_{pA \rightarrow \phi\phi^\dagger} = \sum_{a,b} \int dx_1 dx_2 f_{a/h_1}(x_1) f_{b/h_2}(x_2) \hat{\sigma}_{pA \rightarrow \phi\phi^\dagger}^{(a,b)}, \quad (3.3)$$

where $\hat{\sigma}_{pA \rightarrow \phi\phi^\dagger}^{(a,b)}$ is the partonic cross section for DM production in the scattering of two partons a and b , and f_{a/h_1} (f_{b/h_2}) are the corresponding parton distribution functions of parton a (b) within hadron h_1 (h_2). In our simulation, we employ the leading order PDF set NNPDF2.3LO [62, 63], and we fix the factorization scale to $\mu_F = m_{Z'}$. In the following, we will consider as a benchmark a scalar DM candidate ϕ of mass $m_\phi = 750 \text{ MeV}$ and $U(1)_B$ charge $z_\phi = 3$, while we vary the mass of the force carrier in the range $m_{Z'} \in [2, 7] \text{ GeV}$. In this case, the branching fraction $\text{BR}(Z' \rightarrow \phi\phi^\dagger)$ is of $\mathcal{O}(1)$ [7] and, throughout this work, we assume it to be equal to 1.

Detection of leptophobic DM occurs via scattering of DM particles with the nuclei in the detector. Our search strategy focuses on the deep-inelastic scattering (DIS) signature,




Figure 4. Geometric acceptance for leptophobic scalar DM particles as a function of the detector location relative to the beam axis. The different colored lines correspond to different masses of the Z' mediator. The drop of the acceptance when going towards the on-axis location is a consequence of the scalar nature of the DM candidate.

which is the most relevant one at multi-GeV energies. As for the dark photon model discussed in section 2, we expect better signal/background discrimination power when going off-axis [7]. The number of signal events is computed according to the formula

$$N_{\text{sig}} = \int dz \rho(z) \int_S dS \int dE \frac{d^2 N_\phi(E, S)}{dE dS} \sigma_{\phi(\phi^\dagger)N \rightarrow \phi(\phi^\dagger)N}, \quad (3.4)$$

where z denotes the coordinate along the beam axis, $\rho(z)$ is the number density of nucleons in the detector, S is its surface orthogonal to the z -direction, and $\sigma_{\phi(\phi^\dagger)N \rightarrow \phi(\phi^\dagger)N}$ is the deep-inelastic DM-nucleon scattering cross section. The doubly differential flux $d^2 N_\phi(E, S)/(dE dS)$ of DM particles per unit area dS and per unit energy dE is computed on-the-fly by MADDUMP.

In figure 4 we show the geometric acceptance, which indicates what fraction of DM particles crosses the DUNE-PRISM detectors, as function of the off-axis location of the detectors for different masses $m_{Z'}$ of the force carrier. We observe a distinctive pattern for higher masses: the acceptance steeply increases as the detectors are moved off-axis and then flattens at the largest off-axis angles, θ_{OA} , attainable at DUNE-PRISM. This result is consistent with the ones shown in figure 4 of ref. [7] and is a consequence of our DM candidate being a scalar. Indeed, this behavior can be understood by considering the differential production cross section for scalar DM particles, which reads $d\sigma/d\theta_{\text{OA}} \propto (1 - \cos^2 \theta_{\text{OA}}) \sin \theta_{\text{OA}}$ in the center-of-mass frame. The suppression in the forward direction ($\theta_{\text{OA}} = 0$) is a consequence of angular momentum conservation.

Because of the increasing acceptance at large off-axis angles, combined with the expectation of lower neutrino-induced backgrounds far away from the beam axis, we consider not only the off-axis angles attainable in DUNE-PRISM, but also a hypothetical detector location at $\Delta x_{\text{OA}} = 60 \text{ m}$ ($\theta_{\text{OA}} = 104.15 \text{ mrad}$). This distance is similar to the one of

Experiment	Baseline [m]	Mass [t]	Off-axis [mrad]	Exposure [pot]	n_{bkg}	$\varepsilon_{m_{Z'}=2 \text{ GeV}}$	$\varepsilon_{m_{Z'}=7 \text{ GeV}}$
DUNE on-axis	574	68	0	5.5×10^{21}	2.6×10^7	1.8×10^{-3}	3.5×10^{-5}
DUNE @ 30 m	574	68	52.26	5.5×10^{21}	2.9×10^5	2.5×10^{-3}	1.6×10^{-3}
DUNE @ 60 m	574	68	104.15	5.5×10^{21}	3950	6.3×10^{-4}	1.0×10^{-3}
ICARUS-NuMI	789	480	~ 100	2.5×10^{21}	1600	1.5×10^{-4}	2.7×10^{-4}

Table 1. Summary of experimental setups considered for leptophobic DM. We have used a primary proton energy of 120 GeV and the expected exposure after 5 years for both DUNE and ICARUS-NuMI. A cut on the visible energy ($E_{\text{vis}} > 3 \text{ GeV}$) is applied to all signal and background events, with n_{bkg} being the number of background events that pass the cut. The two rightmost columns indicate the geometric acceptance at two different Z' masses.

the ICARUS-NuMI detector relative to the NuMI beam. According to the investigations carried out in ref. [8], we expect that going that far off-axis leads to close-to-optimal sensitivity as it reduces the background to merely a few thousand events for an exposure of about 2.5×10^{21} pot. While exploiting off-axis distances as large as 60 m in DUNE would require significant civil construction, doing so might be interesting if hints for leptophobic DM interactions should be found in DUNE-PRISM. We use MADDUMP to compute the simulated deep-inelastic neutrino-nucleon scattering events within the detector. For reference, in table 1 we collect the main parameters of different DUNE-PRISM configurations as well as ICARUS-NuMI. In simulating ICARUS-NuMI, we have used the NuMI flux from [64].

3.2 Backgrounds

As DM in the leptophobic model scatters only on hadrons, the main irreducible background channel is neutral current (NC) neutrino scattering. Charged current (CC) interactions might also lead to signal-like signatures when the final state charged lepton is misidentified as a charged pion. We expect the latter contribution to be smaller, as the NC and CC cross sections are of similar size, and hence we neglect it in the following.

The expected background distributions are obtained as described in section 2.2, i.e. by using GENIE to simulate events and then weighting the events with the appropriate on- and off-axis neutrino fluxes. For studies at $\Delta x_{\text{OA}} = 60 \text{ m}$, we compute the neutrino flux based on DUNE’s flux Ntuples published in ref. [50], while for smaller off-axis distances we use directly the flux histograms from the same reference. Considering only data with $E_{\text{vis}} > 3 \text{ GeV}$ allows us to neglect any contributions from resonant scattering, quasi-elastic scattering, and coherent scattering, leaving us with NC deep-inelastic neutrino scattering as the only relevant background process.

The predicted signal and background spectra for the leptophobic model are shown in figure 5. As for the dark photon model (cf. figure 2), we find that backgrounds are dramatically suppressed when going off-axis. NC interactions with a given visible energy E_{vis} are typically caused by neutrinos with an energy $E_\nu \gg E_{\text{vis}}$, implying that they are very sensitive to the high-energy tail of the neutrino spectrum. As the latter is strongly

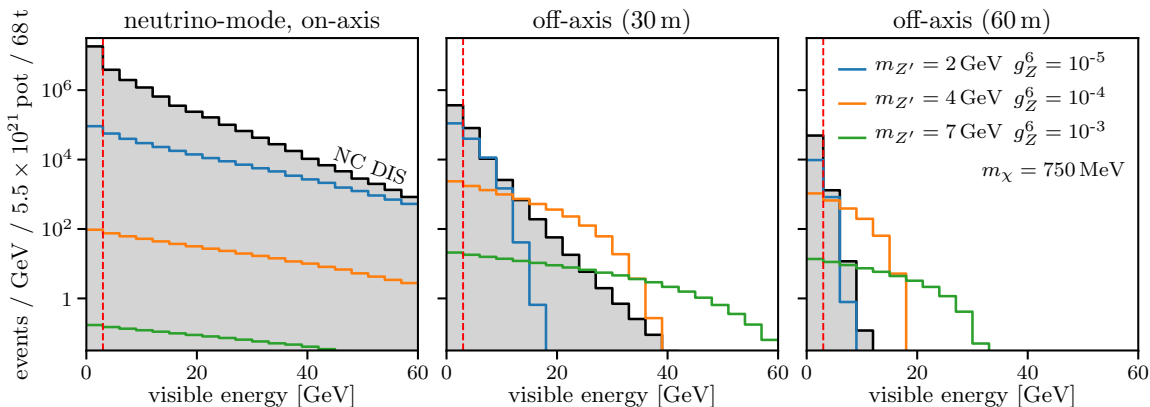


Figure 5. Predicted DM scattering signal in the leptophobic model and corresponding backgrounds for DUNE-PRISM with 5.5×10^{21} pot of neutrino-mode data, corresponding to 5 years of running. Colored histograms show the signal rate for different exemplary model parameter points. The gray shaded background histogram is based on a simulation of NC deep-inelastic neutrino-nucleus scattering. The left panel is for the on-axis location, while the center and right panels correspond to data taken at 30 m (52.26 mrad) off-axis and at a hypothetical detector location 60 m (104.15 mrad) off-axis, respectively. The red dashed line indicates the energy threshold $E_{\text{vis}} > 3$ GeV imposed in our analysis. We observe that backgrounds are significantly suppressed off-axis, especially at high energies. The signal, on the other hand, is even enhanced there for large Z' masses, as explained in section 3.1.

suppressed off-axis, so is the rate of NC background events to our DM search. Because of the lower backgrounds as well as the increased geometric acceptance off-axis (see figure 4), we observe a dramatic increase in the signal-to-background ratio, especially for $m_{Z'}$ in the multi-GeV range.

3.3 Sensitivity to leptophobic dark matter

To estimate the sensitivity of DUNE-PRISM to leptophobic DM, we follow the same statistical procedure as in section 2.4, employing in particular the log-likelihood ratio from eqs. (B.1) and (B.4) which depends on the model parameters $\Theta = (g_Z^6, m_{Z'}, m_\chi)$ in this case. Our results are shown in figure 6, comparing once again a total rates (cut & count) analysis ($n_{\text{bins}} = 1$, left panel) to an analysis utilizing spectral information ($n_{\text{bins}} = 57$ equal-width bins between 3 GeV and 60 GeV, right panel). We find that the spectral analysis clearly outperforms the total rates fit, as can be understood from the event spectra in figure 5. As for leptophilic DM (see section 2.4), the DUNE-PRISM approach plays out its strengths especially for the total rates analysis. Unlike for leptophilic DM, however, on-axis only running is never optimal in the leptophobic model, no matter which type of analysis is used. An off-axis-only run would, however, be competitive with the DUNE-PRISM strategy if a spectral analysis is performed. This behavior can be understood from the better signal acceptance and lower backgrounds that we have discussed above in the context of figures 4 and 5. For the same reason, a hypothetical detector at 60 m off-axis (upper black dotted line in figure 6) would outperform DUNE-PRISM at any of the available on-axis and off-axis locations. Nevertheless, ICARUS-NuMI will do even better using

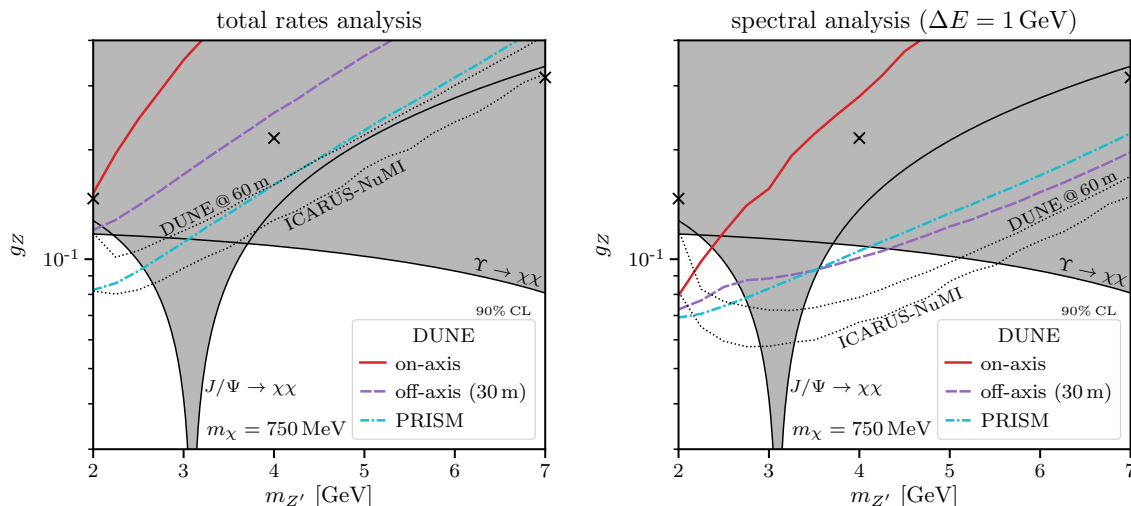


Figure 6. Projected upper exclusion limits for the leptophobic DM model, assuming the DM particle ϕ carries a $U(1)_B$ charge of $z_\phi = 3$. Colored exclusion curves correspond to different DUNE-PRISM running strategies namely on-axis-only running (red solid), off-axis-only running (purple dashed), and equal running times at seven different locations (cyan dot-dashed). Black dotted exclusion curves show the constraints that could be obtained at ICARUS-NuMI using neutrinos from the NuMI beam and at a hypothetical detector placed 60 m off-axis in the DUNE/LBNF beam. Like for the dark photon model (cf. figure 3), we compare a total rates analysis (left panel) to an analysis harnessing the full energy spectrum of events (right panel). We also compare to existing limits from invisible J/ψ and Υ decays at BaBar [65, 66] (gray shaded regions). The black crosses indicate the exemplary model parameter points presented in figure 5.

off-axis neutrinos from the NuMI beam. The excellent performance of ICARUS-NuMI can be attributed to its large mass of 480 t, compared to only 68 t for the combination of the DUNE-PRISM ND-LAr and ND-GAr detector. In addition, the NuMI neutrino flux drops off somewhat faster and becomes softer than the DUNE/LBNF flux when going off-axis. Comparing to existing limits from invisible J/ψ and Υ decays at BaBar [65, 66], we find that DUNE-PRISM will compete with these existing constraints only in small regions around $m_{Z'} \sim 2$ GeV and $m_{Z'} \sim 4$ GeV.

4 Heavy neutral leptons

Heavy neutral leptons (HNLs), often also called sterile neutrinos, are $SU(3)_c \times SU(2)_L \times U(1)_Y$ -singlet fermions whose only (non-gravitational) coupling to the SM is via neutrino mixing. The corresponding operator is

$$\mathcal{L}_{\text{HNL}} \supset y \bar{L} \tilde{H} N, \tag{4.1}$$

where N is the HNL field (a Weyl fermion), L are the left-handed lepton doublets, $\tilde{H} = i\sigma^2 H^*$ is the conjugate of the SM Higgs doublet field, and y is a dimensionless Yukawa coupling. Once the Higgs field acquires a vacuum expectation value, the operator in eq. (4.1) leads to mass mixing between the HNL and the active neutrinos. The HNL thus acquires

the same couplings as the active neutrino and can be produced in any process that produces neutrinos in the SM, unless forbidden by kinematics. Consequently, given a large enough coupling y in eq. (4.1), meson decays in the DUNE target can copiously produce HNLs.

Feynman diagrams involving the mixing operator from eq. (4.1) also admit HNL decays into various final states involving neutrinos, charged leptons, and/or hadrons. If the HNL is sufficiently long-lived (that is, if y is not too large), some of these decays will occur inside the near detector, leaving unique signatures.

The sensitivity of the DUNE near detectors to HNLs has been studied before in refs. [16, 37], albeit for an on-axis configuration only. Current global constraints on HNLs have been compiled for instance in refs. [67–70].

4.1 Production of heavy neutral leptons

In our simulations of HNL production and decay, we closely follow ref. [16], and we use a modified and expanded version of the NUSHOCK code [71] accompanying ref. [16].² More precisely, we predict the HNL flux in DUNE by starting from the simulated neutrino production events released by the DUNE collaboration [50], assuming a total luminosity of 5.5×10^{21} pot, all taken in neutrino mode. Each event in these files, which are based on DUNE’s full Monte Carlo simulation of the target station and decay volume, corresponds to the production of a single SM neutrino. To obtain HNL production events instead, we extract the kinematics of the neutrinos’ parent mesons from the Monte Carlo event files and then use NUSHOCK to re-generate their decays, enforcing the production of an HNL instead of an SM neutrino in the final state. NUSHOCK accounts for the reduced branching ratio to the HNL final state by including an appropriate weight factor for each event.

Unfortunately, the DUNE/LBNF simulation includes only neutrino production in pion and kaon decays, but not in charm decays. This is understandable, given that neutrinos from charm decay are completely irrelevant in the SM. For HNL searches, however, charm decays are very important because at HNL masses larger than the kaon mass, they are the only kinematically allowed HNL production channel. We therefore estimate the flux of charm mesons following once again refs. [16, 71]. We estimate the charm production rate based on the cross section $\sigma_{c\bar{c}} = 2 \mu\text{b}$ for $c\bar{c}$ (open charm) production in proton-proton scattering at the center-of-mass energy $\sqrt{s} = \sqrt{2(1 \text{ GeV}) \cdot (120 \text{ GeV})} \simeq 15 \text{ GeV}$, see figure 16 in ref. [72]. Taking the ratio of $12\sigma_{c\bar{c}}$ to the total proton-¹²C cross section of $\sigma_{pC} = 331.4 \text{ mb}$ [73] gives the rate of charm production per pot. Multiplying further by the fragmentation fraction into D_s , $f_{D_s} = 0.077$ [74], we obtain the rate of D_s production. To generate the D_s kinematics in the center-of-mass frame, we approximate the momentum dependence of the differential D_s production cross section in the center-of-mass frame as [75]

$$\frac{d^2\sigma}{dx dp_{T,D_s}} \propto (1 - |x|)^n \exp(-b p_{T,D_s}^2), \quad (4.2)$$

where p_{T,D_s} is the D_s transverse momentum and $x \equiv 2p_{z,D_s}/\sqrt{s}$, with p_{x,D_s} the D_s momentum along the beam axis. For the numerical coefficients, we use the values $n = 6.1$ and

²As it has been pointed out in [25] there are significant discrepancies in the computation of the HNL branching ratios. However, these discrepancies do not affect our conclusions.




Figure 7. The differential flux of HNLs at the DUNE near detector site (574 m baseline) for different HNL masses (left, middle, and right panels) and for different off-axis angles (different colors and line styles). For illustrative purposes, we have made the unphysical assumption $|U_{e4}|^2 = |U_{\mu4}|^2 = |U_{\tau4}|^2 = 1$. In other words, we have assumed that all neutrino production processes in the DUNE/LBNF beamline are replaced by the corresponding HNL production processes. The fluxes shown here can therefore be linearly rescaled to smaller and more realistic values of $|U_{\alpha4}|^4$. Note the different vertical and horizontal axis scales in the three panels.

$b = 1.08$ [75], which have been measured in the Fermilab E769 experiment using a 250 GeV proton beam impinging on a fixed target.

Our predicted HNL flux as a function of HNL energy is shown in figure 7 for different HNL masses and different off-axis angles. The left panel ($M = 10$ MeV) corresponds to a practically massless HNL that is dominantly produced in pion decays; the HNLs in the middle panel ($M = 200$ MeV) are too heavy to benefit from this production mode, and can only be produced in kaon decays, which explains the much lower flux; at $M = 1$ GeV (right panel), only D meson decays contribute to HNL production. In figure 7, we have assumed $|U_{e4}|^2 = |U_{\mu4}|^2 = |U_{\tau4}|^2 = 1$, implying that all meson decays involving neutrinos in the SM are replaced by the corresponding decays into HNLs.

Not surprisingly, the HNL spectrum becomes softer at larger off-axis angles because higher energy parent mesons are more strongly boosted in the forward direction. This softening of the HNL spectrum when going off-axis is fully analogous to the off-axis softening of the SM neutrino spectrum. Both on-axis and off-axis, the spectra of heavier HNLs (from kaon and charm decays) are significantly harder than the spectra of light HNLs because the smaller boost factors of heavy mesons render their decays more isotropic.

4.2 Heavy neutral lepton decay

To determine the rate of HNL decays inside the DUNE near detectors, we follow the trajectories of all simulated HNLs to determine which of them cross the detectors. We consider both the segmented liquid argon time projection chamber referred to as ArgonCube or ND-LAr by the DUNE collaboration and the pressurized gaseous argon TPC (“multi-purpose detector” or ND-GAr). The ND-LAr detector is box-shaped, with a width (perpendicular to the beam axis) of 7 m, a height of 3 m, and a depth (along the beam axis) of 5 m. The ND-GAr detector is cylindrical, with the cylinder axis oriented horizontally and perpen-

dicular to the beam axis. The detector’s width is 5 m, and the cylinder radius is 2.6 m [23]. We compute the 3D coordinates at which each trajectory enters and exits the detectors by using some elementary geometry, courtesy of the TRIMESH Python package [76]. We then weight each HNL with the probability for decaying inside either of the two detectors.

HNL lifetimes and decay branching ratios into different final states as functions of mass are once again computed using NUSHOCK. In setting limits, we will consider in particular the following final states:

1. $\nu e^+ e^-$ because it has the largest branching ratio among the visible decay channels for HNL masses below the muon threshold. The channel contributes appreciably also at higher masses, and, moreover, backgrounds are small in this channel.
2. $\nu e^\pm \mu^\mp$ and $\nu \mu^+ \mu^-$ due to the large branching ratios for heavy HNLs as well as low background levels
3. $\nu \pi^0$, $e^\pm \pi^\mp$ and $\mu^\pm \pi^\mp$, which dominate at intermediate HNL masses, but are also important at large masses. These channels suffer from large backgrounds that arise from NC or CC neutrino interactions with emission of an extra pion.

For each decay channel, we have used NUSHOCK to simulate a large sample of HNL decays at rest. For each HNL in the DUNE/LBNF beam that decays inside the detector, we randomly pick one of the decay events and boost it from the HNL rest frame to the laboratory frame to obtain the 4-momenta of the observable decay products.

The HNL decay branching ratios are plotted in figure 8 as a function of HNL mass (see also ref. [16]). We see that, at low HNL mass $\lesssim m_\pi$, the dominant decay mode is invisible, $N \rightarrow 3\nu$, followed by $N \rightarrow \nu e^+ e^-$. Above the pion threshold, two-body final states involving charged or neutral pions begin to dominate. Nevertheless, fully leptonic three-body final states ($\nu e^+ e^-$, $\nu \mu^+ \mu^-$, $\nu e^\pm \mu^\mp$) remain relevant, and $N \rightarrow \nu e^\pm \mu^\mp$ even becomes the strongest visible decay mode at $m_N \gtrsim 1.2$ GeV. Decay modes involving kaons, ρ mesons, and other hadrons become important as the corresponding kinematic thresholds are crossed.

4.3 Backgrounds and analysis cuts

To estimate the backgrounds to an HNL search in the DUNE near detectors, we have used GENIE v3.00.06 [49] to simulate CC and NC neutrino interactions on argon. We have simulated 3×10^6 events per flavor, one million each within each of the neutrino energy intervals $[0, 20]$ GeV, $[20, 40]$ GeV, and $[40, 60]$ GeV. Within each interval, neutrino energies are distributed according to the on-axis DUNE flux. To obtain off-axis event samples, events are appropriately re-weighted with the ratio of the off-axis and on-axis fluxes. The rationale for dividing the simulation into three different energy ranges is to obtain more Monte Carlo statistics at high energies, where much of the sensitivity to HNLs is coming from.

We process the simulated events through the background simulation implemented in NUSHOCK [16, 71]. This involves a very simple detector simulation that applies Gaussian




Figure 8. Branching ratios of HNL decay modes as a function of mass. We have assumed flavor-universal mixing, i.e. $|U_{e4}| = |U_{\mu 4}| = |U_{\tau 4}|$.

smearing to the momenta of the final state particles and uses a kinetic energy threshold to determine which of them are reconstructed. It rejects any events containing hadrons other than pions or atomic nuclei, thus exploiting the fact that the HNL decay modes that we consider do not contain heavy hadrons, while many potential background processes are from deep-inelastic neutrino scattering events that typically involve a lot of hadronic activity. Note that π^0 s, which are by default not decayed in GENIE, are instead decayed in NUSHOCK and the analysis is carried out on the two final state photons. Charged pions, on the other hand, are retained undecayed due to their much longer lifetime.

The next step consists of simple particle misidentification rules. In particular:

- Pions are misreconstructed as muons if their randomly chosen track length is sufficiently long (> 2 m).
- e^+e^- pairs are misreconstructed as (converted) photons if their angular separation is below a threshold (3°).
- Photons that convert after less than 2 cm are reconstructed as electrons or positrons.

The list of final state particles after threshold cuts and misidentification is then compared to the list of final state particles expected for the signal channel under consideration. Only if the number and type of each particle match between the signal and background, the event is retained (exclusive analysis). We do, however, conservatively assume no charge identification capabilities even though the ND-GAr detector will have a magnetic field and

should therefore be able to efficiently distinguish positively and negatively charged final state particles.

To further suppress backgrounds, we exploit the fact that HNL decay products are typically strongly boosted in the forward direction, while background events have a more isotropic topology. We implement a cut on the angle θ between the mean direction of the two visible would-be HNL decay products and the beam axis. For given HNL mass M , we set the threshold at $\theta < M/(E_1 + E_2)$, where E_1 and E_2 are the energies of the two would-be HNL decay products. This threshold value is an estimate for the forward boost of the decay products in a real HNL decay.

We find the dominant background contributions for the different HNL decay channels to be as follows:

1. For $N \rightarrow \nu e^+ e^-$, by far the most important background is due to misidentified photons.
2. $N \rightarrow \nu \mu^+ \mu^-$ is most easily mimicked by CC ν_μ interactions with one real muon and one misidentified charged pion.
3. Similarly, $N \rightarrow \nu e^\pm \mu^\mp$ suffers from a background due to CC ν_μ interactions with a real muon and a photon that is misidentified as an electron.
4. Backgrounds to $N \rightarrow e^\pm \pi^\mp$ arise mostly from CC ν_e interactions with pion production. A smaller contribution comes from NC neutrino interactions with a real pion and a photon misidentified as an electron.
5. Similarly, $N \rightarrow \mu^\pm \pi^\mp$ is affected by a large background due to CC ν_μ interactions with pion production.
6. Finally, in $N \rightarrow \nu \pi^0$, the source of the large background is NC neutrino interactions with pion production.

We have also considered neutrino trident production (which is not simulated by GENIE) as a possible source of background. However, based on the calculation in ref. [77], we conclude that this background will be negligible.

In figure 9, we compare the signal and background predictions for several HNL masses and decay channels, both on-axis (blue) and off-axis (purple). We have chosen the same representative HNL masses as in figure 7, and the decay channel shown for each of them is the most sensitive one at this particular mass (in the absence of backgrounds). The values of the mixing matrix elements $|U_{\alpha 4}|^2$ (assumed to be the same for all flavors $\alpha = e, \mu, \tau$) are chosen the 95% CL limits at the given masses. All channels shown here involve two visible final state particles, and we show the spectra for both of them separately. (In the case of $N \rightarrow \nu \pi^0$, the two final state particles are the two photons from π^0 decay.) If the two final state particles are identical (or identical up to their charge, which we assume is not measured), the upper plot shows the spectrum of the harder of the two, while the lower plot is for the softer one. We observe that, for all channels, background levels at high energies are lowered by several orders of magnitude when going off-axis. The signal spectra,




Figure 9. Signal and background event spectra at the DUNE near detector complex for specific HNL decay modes. Blue and purple unshaded histograms correspond to the number of signal events per GeV at the on-axis location and at 18 m off-axis, respectively. Shaded histograms represent the corresponding backgrounds. The three panels correspond to the most sensitive decay channels in the absence of backgrounds at $M = 10$ MeV (left panel), $M = 200$ MeV (middle panel), and $M = 1000$ MeV (right panel), respectively, as indicated in the plots. For all decay channels included here ($N \rightarrow \nu e^+ e^-$ in the left panel, $N \rightarrow e^\pm \pi^\mp$ in the middle panel, $N \rightarrow \nu + (\pi^0 \rightarrow \gamma\gamma)$ in the right panel), we show the spectra of the two visible final state particles separately. If the particles are identical (up to their charge, which we conservatively assume not to be measured), the upper histogram corresponds to the more energetic one, while the lower histogram is for the softer particle. Note that the sensitivity estimates discussed below are based on the full two-dimensional event distributions. We observe that the signal-to-background ratio improves significantly when going off-axis.

on the other hand, drop somewhat more slowly. This is true especially at low energies, where the HNLs’ forward boost is relatively small and so their angular distribution is more isotropic than the one of SM neutrinos, which are responsible for the background.

4.4 Sensitivity to heavy neutral leptons

We are now ready to estimate the sensitivity of the DUNE near detectors to HNL decays. We use again the log-likelihood function from eq. (B.1), but here we do so separately for ND-GAr and ND-LAr. This way, our analysis benefits from the much better signal-to-background ratio in ND-GAr which comes from the fact that the HNL signal scales with the detector volume, while the background scales with its mass. The two independent χ^2 values are only added up in the very end. As the HNL decay final states that we consider here contain two visible particles, we bin the events in two dimensions, corresponding to the energies of the two particles. This turns out to be very important for optimizing the sensitivity. One problem with working with two-dimensional histograms is that it is difficult to generate sufficient background Monte Carlo statistics in all bins. This can be problematic because bins that contain zero background events simply due to limited simulation statistics can lead to sensitivity estimates that are too optimistic. To mitigate this problem, we choose larger bins at higher energies: our bin width is 1 GeV for particle energies below 10 GeV, 2 GeV for particle energies up to 20 GeV, and 4 GeV for particle energies up to

40 GeV. To be on the safe side, we also identify bins in which the background prediction is exactly zero. We then consider averages of the background rate over increasingly larger neighborhoods of each such bin until we obtain a non-zero rate. If that average rate is above 0.1 events, we exclude the problematic bin from our analysis.

As systematic uncertainties, we include a 10% normalization error, which we assume to be uncorrelated between different HNL decay channels and between the signal and background. These uncertainties are described in terms of nuisance parameters with Gaussian priors. To obtain the final sensitivity, our χ^2 function is minimized over all nuisance parameters.

Our results are shown in figures 10 and 11. The four panels in figure 10 correspond to different assumption on the HNL couplings (coupling to ν_e only, to ν_μ only, to ν_τ only, and to all three active neutrino flavors universally). We observe rich structure in these plots: first, we see that the projected exclusion regions have the wedge shape that is typical for long-lived particle searches. At too small mixing, HNL production is suppressed beyond the detectable level and HNLs are so long-lived that they mostly decay far beyond the detector. At too large mixing, HNLs are abundantly produced, but most of them decay before reaching the detector. Away from kinematic thresholds, the upper limits on the mixing matrix elements $|U_{e4}|^2$, $|U_{\mu4}|^2$, $|U_{\tau4}|^2$ (bottom edges of the wedge-shaped regions in figure 10) scale roughly as $|U_{\alpha4}|_{\text{limit}}^2 \propto M^{-3}$. This can be understood as follows: the HNL production rate scales as $|U_{\alpha4}|^2 M^2$ over wide mass ranges, with the proportionality to M^2 reflecting the chiral suppression that otherwise affects many leptonic meson decays. The HNL decay rate in the HNL rest frame scales as $|U_{\alpha4}|^2 M^5$, as can be seen from dimensional analysis. In the laboratory frame, this scaling changes to $|U_{\alpha4}|^2 M^6$ due to relativistic time dilation. Simultaneously, the opening angle of the HNL beam coming from its Lorentz boost grows with M , implying that the fraction of HNLs crossing the detector drops as M^{-2} . Overall, these arguments show that the experimental count rate scales as $|U_{\alpha4}|^4 M^6$, with deviations being observed close to kinematic thresholds. Moreover, the geometric scaling factor is not always exactly M^{-2} , depending on how exactly the flux in a given channel drops with angle.

Let us now discuss the kinematic thresholds visible in figure 10. For instance, for N - ν_e coupling only (top left panel of figure 10), we see that, at $M \sim m_\pi \sim 140$ MeV, HNL production in pion decays becomes kinematically forbidden, leading to a dent in the sensitivity in the $N \rightarrow \nu e^+ e^-$ channel. Simultaneously, however, new HNL decay modes ($N \rightarrow e^\pm \pi^\mp$ and $N \rightarrow \nu \pi^0$) become allowed, compensating to some extent for this loss of sensitivity due to their large branching ratios. The next kinematic threshold occurs at $M \sim m_K \sim 490$ MeV, when also HNL production in kaon decays becomes forbidden. Beyond this threshold, only HNLs from charm decay contribute, but since charm production in DUNE/LBNF is relative inefficient due to the low primary proton energy of 120 GeV, this leads to a significant drop in sensitivity. Nevertheless, the DUNE near detectors will be able to probe a large chunk of parameter space up to $M \sim m_{D_s} \sim 2$ GeV, where also HNL production in charm decays becomes forbidden.

For HNL couplings to ν_μ (top right panel of figure 10), all thresholds are shifted by about $m_\mu \sim 100$ GeV because each HNL needs to be produced together with a muon. For




Figure 10. Sensitivity to HNLs in the DUNE near detectors ND-LAr and ND-GAr as function of HNL mass. Solid lines correspond to the DUNE-PRISM running strategy with data taking both on-axis and at six different off-axis locations (6 m, 12 m, 18 m, 24 m, 30 m, 36 m), with equal exposure for each location. The total number of protons on target assumed here is 5.5×10^{21} , corresponding to 5 years of running in neutrino mode. Dotted lines show the sensitivity for on-axis data taking only, and thin dashed lines have been computed from a DUNE-PRISM analysis that neglects backgrounds. The four panels correspond to different HNL coupling structures, namely mixing with ν_e only (top left panel), with ν_μ only (top right panel), with ν_τ only (bottom left panel), and flavor-universal mixing $|U_{e4}|^2 = |U_{\mu4}|^2 = |U_{\tau4}|^2 \equiv |U|^2$ (bottom right panel).

HNL couplings exclusively to ν_τ (bottom left panel of figure 10), production thresholds do not play as important a role because charm decays are the only production mode at all HNL masses due to the requirement of producing a τ alongside the HNL. The small kink visible at $M \sim 190$ MeV can be understood from the D_s - τ mass difference. Beyond the kink, HNL production in $D_s \rightarrow N + \tau$ decays is kinematically forbidden, leaving only the off-shell decays $D_s \rightarrow \nu_\tau + (\tau^* \rightarrow N e \nu_e, N \mu \nu_\mu, N + \text{hadrons})$ as viable HNL production modes.

Comparing on-axis and off-axis sensitivities, we notice that going off-axis does not lead to significant benefits, in spite of the much better signal-to-background ratio. The reason is that, with the cuts discussed above, background suppression is fairly effective

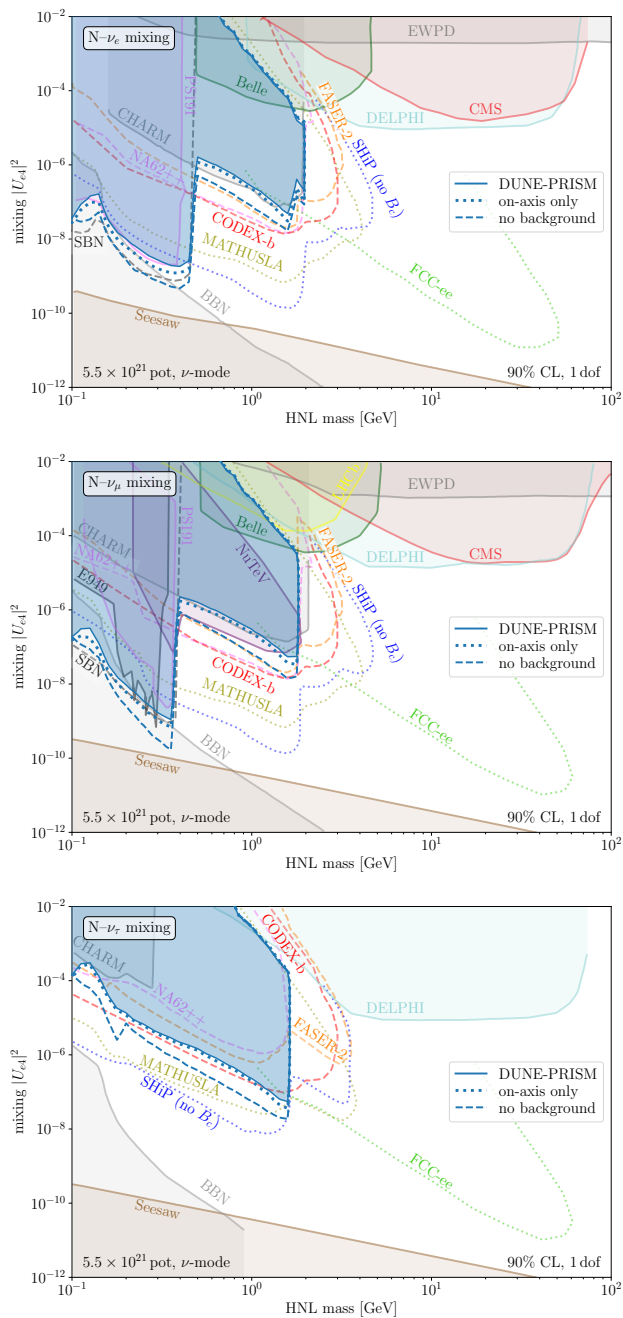


Figure 11. The DUNE near detectors compared to other experiments sensitive to HNLs. The top, middle, and bottom panels show the sensitivity to the squared mixing matrix elements $|U_{e4}|^2$, $|U_{\mu 4}|^2$, and $|U_{\tau 4}|^2$, respectively, assuming that only one of them is non-zero at a time. In this plot, the decay channels νe^+e^- , $\nu\mu^+\mu^-$, $\nu e^\pm\mu^\mp$, $\nu\pi^0$, $\pi^\pm e^\mp$, $\pi^\pm\mu^\mp$ are combined. As in figure 10, solid blue curves show the sensitivity of DUNE-PRISM (5.5×10^{21} pot in neutrino mode, equally split between the on-axis position and six different off-axis locations). Dotted blue contours show results for on-axis running only, and thin dashed blue curves represent a hypothetical background-free analysis. The contours shown in the background correspond to existing limits (filled) and sensitivities of planned experiments (unfilled dashed and dotted). These limits are taken from the compilation in ref. [54] and from ref. [78].

even on-axis. It is important to note, though, that going off-axis also does not significantly harm the sensitivity for any channel, in spite of the lower off-axis fluxes. This means that the search for HNLs can be carried out truly parasitically to DUNE’s main oscillation program. No matter where the various near detectors are placed at any given moment, they will contribute in a useful way to the HNL sensitivity. Notably, the SAND (System for On-Axis Neutrino Detection) beam monitor, which will always remain on-axis and has not been included in our estimates, can be harnessed for the HNL search as well.

In figure 11, we compare DUNE’s sensitivity to an array of current and future HNL constraints, combining all six HNL decay channels analyzed above. We do so for the same three scenarios as in figure 10: on-axis running only, the realistic DUNE-PRISM running strategy with equal amounts of data collected at seven different locations relative to the beam axis (0 m, 6 m, 12 m, 18 m, 24 m, 30 m, 36 m [50]), and a hypothetical background-free DUNE-PRISM search. The comparison reveals that DUNE will be able to somewhat improve on existing limits for HNL couplings to ν_e and ν_μ , and will go far beyond what is currently possible for HNL couplings to ν_τ . (While this conclusion hinges to some extent on our treatment of charm production in LBNF, we are confident that, given the very conservative choice we have made for the charm production cross section, the sensitivity of the real experiment will likely be better than our estimate.)

Comparing to planned experiments that may happen on a timescale similar to DUNE, we note that the ones that are most competitive with DUNE at low HNL masses (below the kaon threshold) are the Fermilab short-baseline experiments (“SBN”), that is MicroBooNE, SBND, and ICARUS. Other planned experiments (FASER-2, CODEX-b, NA62++, MATHUSLA, SHiP, FCC-ee) probe higher HNL masses than DUNE because they operate at higher beam energies. They are, however, not able to reach the luminosities achievable in DUNE.

We should also keep in mind that the limits shown here are conservative in several ways, in particular they are based on only 5 years of data taking, and we have assumed no charge identification capabilities.

5 Summary and conclusions

In this work, we have explored the benefits of the DUNE-PRISM detectors for exploring physics beyond the SM, namely two models of light DM and a scenario featuring heavy neutral leptons (or sterile neutrinos). We have focused in particular on the capability of these detectors to move in and out of the beam axis and have found that exploiting this capability is highly beneficial in some scenarios, and never disadvantageous.

For MeV-GeV-scale scalar DM coupled via a dark photon with kinetic mixing, we have found that a detector placed off-axis will see a significant reduction in backgrounds from SM neutrino scattering. While also the signal due to DM-electron scattering is suppressed, the signal-to-background ratio is improved significantly. This implies that a realistic DUNE-PRISM running strategy that combines on-axis and off-axis measurements will be ideal for constraining such a scenario. This is especially true if a simple cut & count analysis is performed; for a full spectral analysis, an on-axis-only measurement would perform equally

well as DUNE-PRISM. Overall, a 5-year run of DUNE-PRISM (5.5×10^{21} pot) will be able to improve existing limits on dark photon-mediated light DM by up to a factor of a few, as shown in figure 3.

If light scalar DM couples through a leptophobic Z' gauge boson rather than a dark photon, we are in the interesting situation that, for some parameter ranges (namely those with relatively heavy Z'), the DM flux increases away from the beam axis, while the background from SM neutral current neutrino interactions falls off. Consequently, the DUNE-PRISM strategy is always advantageous compared to an on-axis-only run, see figure 6. After 5 years, DUNE-PRISM limits will be competitive with existing ones, and better in some parameter regions.

Turning finally to heavy neutral leptons (figure 11), we find once again that taking data both on-axis and off-axis as in DUNE-PRISM never hurts the sensitivity. For some decay channels, especially those with large backgrounds like $N \rightarrow \nu\pi^0$, off-axis running is highly beneficial. For channels where backgrounds are lower or can be well distinguished from the signal by using spectral information, the sensitivity is similar off-axis and on-axis. Compared to existing limits, DUNE-PRISM after 5 years will be competitive and in some parameter regions better than existing limits for heavy neutral lepton mixing with ν_e and ν_μ . Significant new territory will be covered for couplings to ν_τ thanks to charm production in the DUNE/LBNF target.

We conclude that the DUNE-PRISM detectors are a versatile new tool for probing numerous extensions of the SM. Their unique capability of moving off-axis will only improve the sensitivity to such scenarios, implying that a rich program of new physics searches can be carried out concurrently with DUNE's neutrino oscillation program without requiring an adaptation of the running strategy.

Acknowledgments

We are deeply grateful to Tommaso Boschi for answering numerous questions about his NUSHOCK code [71], to Laura Fields for useful discussions about the DUNE fluxes, and to Pedro Machado for helping us in the comparison of our results to the results of ref. [13]. We have benefited from useful discussions with Gaia Lanfranchi, Albert de Roeck, and the members of the DUNE-HNL working group. This work has been partly funded by the German Research Foundation (DFG) in the framework of the PRISMA+ Cluster of Excellence and by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 637506, “ ν Directions”). The research of L. B. is supported in part by the Swiss National Foundation under Contracts No. 200020_188464 and No. IZSAZ2_173357.

A Simulation of dark matter production in the dark photon model

In the context of the dark photon model introduced in section 2, we have compared our signal and background predictions with those of ref. [13]. We have found that our estimate of the background due to neutrino-electron scattering agrees well with the results of ref. [13]




Figure 12. Same as figure 1, but including the result of ref. [13] in solid green, and our own *primary-only (HardQCD)* sample in dotted orange.

in both the neutrino- and anti-neutrino mode. In the following, we focus on the estimate of the number of signal events.

As discussed in section 2.1 a crucial aspect in the simulation of the signal events is the modeling of the light meson spectra from the DUNE/LBNF target. In ref. [13], only the production of mesons in the primary proton interaction was considered, and PYTHIA was used as the main simulation tool. In figure 12 we compare our results with those of figure 3 of ref. [13] using the same benchmark point $m_{A'} = 90 \text{ MeV} = 3m_\phi$, $\epsilon^4\alpha_D = 10^{-15}$. Contrary to our expectation, we observe that our *primary-only (SoftQCD)* curve fails to accurately reproduce the results of that paper (solid green line) at any value of Δx_{OA} . As anticipated in section 2.1, the source of the discrepancy is that the authors of ref. [13] have used a different configuration for their PYTHIA simulation. More precisely, they activated the `HardQCD:All` flag, and if we do the same (dotted orange curve in figure 12), we are indeed able to reproduce their results.

We remark that using the `SoftQCD:All` flag is preferable for modeling the proton-proton primary interactions for the case of relatively low energy fixed target experiments such as DUNE/LBNF. (This was shown in ref. [37], which appeared after ref. [13].) In fact, *primary-only (HardQCD)* events tend to be much softer and have a larger angular spread than what is expected from primary proton-proton interactions, as also show in figure 13. Hence, the *primary-only (HardQCD)* sample underestimates the number of signal events on-axis, while being coincidentally similar to the *beam-dump* sample when going off-axis.




Figure 13. Energy spectrum of π^0 s produced in the beam dump in three angular windows: $0 < \theta_{OA} < 10$ mrad (left panel), $20 \text{ mrad} < \theta_{OA} < 30$ mrad (middle panel) and $50 \text{ mrad} < \theta_{OA} < 60$ mrad (right panel). The three curves correspond to the three meson samples introduced in the main text: *primary-only (SoftQCD)* in red dot-dashed, *beam-dump primary-only (HardQCD)* in orange dotted.

B Statistical analysis

In the following, we describe the methodology used to derive sensitivity limits from our predicted signal and background rates in the analyses of sections 2 to 4 in more detail. As stated in these sections we use standard frequentist techniques to test the signal + background hypothesis against simulated background-only data. We employ a Poissonian log-likelihood function, i.e.

$$\begin{aligned}
 -2 \log \mathcal{L}(\Theta, \mathbf{X}) \equiv & 2 \sum_{j=1}^{n_{\text{pos}}} \sum_{i=1}^{n_{\text{bins}}} \left[B_{ij}(\mathbf{X}) + S_{ij}(\Theta, \mathbf{X}) - B_{ij}(0) + B_{ij}(0) \log \left(\frac{B_{ij}(0)}{B_{ij}(\mathbf{X}) + S_{ij}(\Theta, \mathbf{X})} \right) \right] \\
 & + \sum_{c=S,B} \left\{ \left(\frac{X_{c,\text{correl}}^{\text{norm}}}{\sigma_{\text{correl}}} \right)^2 + \left(\frac{X_{c,\text{correl}}^{\text{tilt}}}{\sigma_{\text{correl}}} \right)^2 + \sum_{j=1}^{n_{\text{pos}}} \left[\left(\frac{X_{c,j}^{\text{norm}}}{\sigma_{\text{pos}}} \right)^2 + \left(\frac{X_{c,j}^{\text{tilt}}}{\sigma_{\text{pos}}} \right)^2 \right] \right\}.
 \end{aligned} \tag{B.1}$$

The signal and background rates in the i -th energy bin at the j -th off-axis position including systematic biases, $S_{ij}(\Theta, \mathbf{X})$ and $B_{ij}(\mathbf{X})$, are defined in terms of the rates without biases, $S_{ij}(\Theta, 0)$ and $B_{ij}(0)$, according to

$$\begin{aligned}
 S_{ij}(\Theta, \mathbf{X}) &= [1 + X_{S,\text{correl}}^{\text{norm}}] [1 + X_{S,j}^{\text{norm}}] [1 + X_{S,\text{correl}}^{\text{tilt}}(-1, \dots, 1)_i] [1 + X_{S,j}^{\text{tilt}}(-1, \dots, 1)_i] S_{ij}(\Theta, 0), \\
 B_{ij}(\mathbf{X}) &= [1 + X_{B,\text{correl}}^{\text{norm}}] [1 + X_{B,j}^{\text{norm}}] [1 + X_{B,\text{correl}}^{\text{tilt}}(-1, \dots, 1)_i] [1 + X_{B,j}^{\text{tilt}}(-1, \dots, 1)_i] B_{ij}(0).
 \end{aligned} \tag{B.2}$$

We collectively denote the vector of physical model parameters Θ and the vector of nuisance parameters \mathbf{X} . The nuisance parameters $X_{S,\text{corr}}^{\text{norm}}$ and $X_{S,j}^{\text{norm}}$ describe systematic normalization uncertainties in the signal, while the parameters $X_{S,\text{correl}}^{\text{tilt}}$ and $X_{S,j}^{\text{tilt}}$ parameterize spectral “tilts”: their effect is to pivot the spectrum about its midpoint. The meaning of the corresponding parameters affecting the background ($X_{B,\text{corr}}^{\text{norm}}$, $X_{B,j}^{\text{norm}}$, $X_{B,\text{correl}}^{\text{tilt}}$ and $X_{B,j}^{\text{tilt}}$) is

analogous. (Note that tilt errors are not considered in the HNL analysis of section 4.) All systematic errors are treated as Gaussian and are constrained by the pull terms in the second row of eq. (B.1). As systematic normalization and tilt uncertainties we assume $\sigma_{\text{pos}} = 1\%$ (uncorrelated between different on-/off-axis positions) and $\sigma_{\text{correl}} = 10\%$ (correlated among all positions). The sums in the first line of eq. (B.1) run over $n_{\text{bins}} = 80$ ($n_{\text{bins}} = 57$) energy bins, equally spaced in the interval $[0, 20]$ GeV ($[3, 60]$ GeV) in case of the dark photon model (leptophobic model), and over $n_{\text{pos}} = 7$ on- and off-axis positions (0 m, 6 m/10.45 mrad, 12 m/20.90 mrad, 18 m/31.36 mrad, 24 m/41.81 mrad, 30 m/52.26 mrad, 36 m/62.72 mrad). For comparison we also present results for an on-axis-only run, or for data taken at a fixed off-axis location. In this case, of course, we set $n_{\text{pos}} = 1$. For HNL searches, we use the two-dimensional binning described in section 4.4.

The statistical significance at which a given parameter point Θ is excluded can be estimated from the log-likelihood ratio [79]

$$Z(\Theta) \equiv -2 \log \left(\frac{\mathcal{L}(\Theta, \hat{\mathbf{X}})}{\mathcal{L}(\hat{\Theta}, \hat{\mathbf{X}})} \right), \tag{B.3}$$

where $(\hat{\Theta}, \hat{\mathbf{X}})$ is the combination of model parameters and nuisance parameters that maximizes the likelihood (minimizes χ^2), and $\hat{\mathbf{X}}$ are the nuisance parameters that maximize the likelihood for fixed Θ . $Z(\Theta)$ follows a χ^2 distribution, with the number of degrees of freedom equal to the dimension of Θ . Consequently, the $n\sigma$ exclusion region can be estimated by comparing $Z(\Theta)$ to the $n\sigma$ quantile of that χ^2 distribution.

However, experimental collaboration often present their results as constraints on the “signal strength” μ , which is defined as an overall, energy-independent rescaling factor of the event rate relative to some reference point. In the dark photon model, the signal strength can be taken as the product of coupling constants $\mu \equiv \epsilon^4 \alpha_D$, while in the leptophobic model it is $\mu \equiv g_Z^6$. In this way of analyzing the data, $Z(\Theta)$ is replaced by a one-parameter function

$$Z(\mu) \equiv -2 \log \left(\frac{\mathcal{L}(\mu, \hat{\mathbf{X}})}{\mathcal{L}(\hat{\mu}, \hat{\mathbf{X}})} \right), \tag{B.4}$$

whose values are then compared to the χ^2 distribution with one degree of freedom. In eq. (B.4), $(\hat{\mu}, \hat{\mathbf{X}})$ is once again the parameter point in (μ, \mathbf{X}) at which the likelihood is maximal. In a signal strength analysis, however, only μ is varied and all other model parameters are kept fixed. The statistical question answered by a signal strength analysis is “What is the constraint on the signal strength, assuming we already know all other model parameters.” A multi-parameter fit, on the other hand, asks “What are the preferred parameter regions, assuming the chosen model is the correct one, but we do not have any prior information on any of its parameters.” To make our results comparable to those shown in the literature, we use signal strength analyses throughout.

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

References

- [1] B. Batell, M. Pospelov and A. Ritz, *Exploring Portals to a Hidden Sector Through Fixed Targets*, *Phys. Rev. D* **80** (2009) 095024 [[arXiv:0906.5614](https://arxiv.org/abs/0906.5614)] [[INSPIRE](#)].
- [2] P. deNiverville, M. Pospelov and A. Ritz, *Observing a light dark matter beam with neutrino experiments*, *Phys. Rev. D* **84** (2011) 075020 [[arXiv:1107.4580](https://arxiv.org/abs/1107.4580)] [[INSPIRE](#)].
- [3] P. deNiverville, D. McKeen and A. Ritz, *Signatures of sub-GeV dark matter beams at neutrino experiments*, *Phys. Rev. D* **86** (2012) 035022 [[arXiv:1205.3499](https://arxiv.org/abs/1205.3499)] [[INSPIRE](#)].
- [4] MINIBOONE collaboration, *Low Mass WIMP Searches with a Neutrino Experiment: A Proposal for Further MiniBooNE Running*, [arXiv:1211.2258](https://arxiv.org/abs/1211.2258) [[INSPIRE](#)].
- [5] B. Batell, P. deNiverville, D. McKeen, M. Pospelov and A. Ritz, *Leptophobic Dark Matter at Neutrino Factories*, *Phys. Rev. D* **90** (2014) 115014 [[arXiv:1405.7049](https://arxiv.org/abs/1405.7049)] [[INSPIRE](#)].
- [6] D.E. Soper, M. Spannowsky, C.J. Wallace and T.M.P. Tait, *Scattering of Dark Particles with Light Mediators*, *Phys. Rev. D* **90** (2014) 115005 [[arXiv:1407.2623](https://arxiv.org/abs/1407.2623)] [[INSPIRE](#)].
- [7] B.A. Dobrescu and C. Frugiuele, *GeV-Scale Dark Matter: Production at the Main Injector*, *JHEP* **02** (2015) 019 [[arXiv:1410.1566](https://arxiv.org/abs/1410.1566)] [[INSPIRE](#)].
- [8] P. Coloma, B.A. Dobrescu, C. Frugiuele and R. Harnik, *Dark matter beams at LBNF*, *JHEP* **04** (2016) 047 [[arXiv:1512.03852](https://arxiv.org/abs/1512.03852)] [[INSPIRE](#)].
- [9] C. Frugiuele, *Probing sub-GeV dark sectors via high energy proton beams at LBNF/DUNE and MiniBooNE*, *Phys. Rev. D* **96** (2017) 015029 [[arXiv:1701.05464](https://arxiv.org/abs/1701.05464)] [[INSPIRE](#)].
- [10] P. deNiverville and C. Frugiuele, *Hunting sub-GeV dark matter with the NOvA near detector*, *Phys. Rev. D* **99** (2019) 051701 [[arXiv:1807.06501](https://arxiv.org/abs/1807.06501)] [[INSPIRE](#)].
- [11] G. Magill, R. Plestid, M. Pospelov and Y.-D. Tsai, *Millicharged particles in neutrino experiments*, *Phys. Rev. Lett.* **122** (2019) 071801 [[arXiv:1806.03310](https://arxiv.org/abs/1806.03310)] [[INSPIRE](#)].
- [12] A. de Gouvêa, P.J. Fox, R. Harnik, K.J. Kelly and Y. Zhang, *Dark Tridents at Off-Axis Liquid Argon Neutrino Detectors*, *JHEP* **01** (2019) 001 [[arXiv:1809.06388](https://arxiv.org/abs/1809.06388)] [[INSPIRE](#)].
- [13] V. De Romeri, K.J. Kelly and P.A.N. Machado, *Hunting On- and Off-Axis for Light Dark Matter with DUNE-PRISM*, *Phys. Rev. D* **100** (2019) 095010 [[arXiv:1903.10505](https://arxiv.org/abs/1903.10505)] [[INSPIRE](#)].
- [14] L. Buonocore, C. Frugiuele and P. deNiverville, *Hunt for sub-GeV dark matter at neutrino facilities: A survey of past and present experiments*, *Phys. Rev. D* **102** (2020) 035006 [[arXiv:1912.09346](https://arxiv.org/abs/1912.09346)] [[INSPIRE](#)].
- [15] B. Batell, J. Berger and A. Ismail, *Probing the Higgs Portal at the Fermilab Short-Baseline Neutrino Experiments*, *Phys. Rev. D* **100** (2019) 115039 [[arXiv:1909.11670](https://arxiv.org/abs/1909.11670)] [[INSPIRE](#)].
- [16] P. Ballett, T. Boschi and S. Pascoli, *Heavy Neutral Leptons from low-scale seesaws at the DUNE Near Detector*, *JHEP* **03** (2020) 111 [[arXiv:1905.00284](https://arxiv.org/abs/1905.00284)] [[INSPIRE](#)].
- [17] M. Battaglieri et al., *US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report*, in *U.S. Cosmic Visions: New Ideas in Dark Matter*, (2017) [[arXiv:1707.04591](https://arxiv.org/abs/1707.04591)] [[INSPIRE](#)].

- [18] SENSEI collaboration, *SENSEI: First Direct-Detection Constraints on sub-GeV Dark Matter from a Surface Run*, *Phys. Rev. Lett.* **121** (2018) 061803 [[arXiv:1804.00088](#)] [[INSPIRE](#)].
- [19] MINIBOONE collaboration, *Dark Matter Search in a Proton Beam Dump with MiniBooNE*, *Phys. Rev. Lett.* **118** (2017) 221803 [[arXiv:1702.02688](#)] [[INSPIRE](#)].
- [20] MINIBOONE DM collaboration, *Dark Matter Search in Nucleon, Pion, and Electron Channels from a Proton Beam Dump with MiniBooNE*, *Phys. Rev. D* **98** (2018) 112004 [[arXiv:1807.06137](#)] [[INSPIRE](#)].
- [21] MICROBOONE collaboration, *Search for Heavy Neutral Leptons Decaying into Muon-Pion Pairs in the MicroBooNE Detector*, *Phys. Rev. D* **101** (2020) 052001 [[arXiv:1911.10545](#)] [[INSPIRE](#)].
- [22] MICROBOONE collaboration, *Search for a Higgs Portal Scalar Decaying to Electron-Positron Pairs in the MicroBooNE Detector*, *Phys. Rev. Lett.* **127** (2021) 151803 [[arXiv:2106.00568](#)] [[INSPIRE](#)].
- [23] DUNE collaboration, *Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report*, *Instruments* **5** (2021) 31 [[arXiv:2103.13910](#)] [[INSPIRE](#)].
- [24] I. Krasnov, *DUNE prospects in the search for sterile neutrinos*, *Phys. Rev. D* **100** (2019) 075023 [[arXiv:1902.06099](#)] [[INSPIRE](#)].
- [25] P. Coloma, E. Fernández-Martínez, M. González-López, J. Hernández-García and Z. Pavlovic, *GeV-scale neutrinos: interactions with mesons and DUNE sensitivity*, *Eur. Phys. J. C* **81** (2021) 78 [[arXiv:2007.03701](#)] [[INSPIRE](#)].
- [26] E. Izaguirre, G. Krnjaic, P. Schuster and N. Toro, *Analyzing the Discovery Potential for Light Dark Matter*, *Phys. Rev. Lett.* **115** (2015) 251301 [[arXiv:1505.00011](#)] [[INSPIRE](#)].
- [27] T. Lin, H.-B. Yu and K.M. Zurek, *On Symmetric and Asymmetric Light Dark Matter*, *Phys. Rev. D* **85** (2012) 063503 [[arXiv:1111.0293](#)] [[INSPIRE](#)].
- [28] PLANCK collaboration, *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13 [[arXiv:1502.01589](#)] [[INSPIRE](#)].
- [29] T.R. Slatyer, *Indirect dark matter signatures in the cosmic dark ages. I. Generalizing the bound on s-wave dark matter annihilation from Planck results*, *Phys. Rev. D* **93** (2016) 023527 [[arXiv:1506.03811](#)] [[INSPIRE](#)].
- [30] T.R. Slatyer, *Indirect Dark Matter Signatures in the Cosmic Dark Ages II. Ionization, Heating and Photon Production from Arbitrary Energy Injections*, *Phys. Rev. D* **93** (2016) 023521 [[arXiv:1506.03812](#)] [[INSPIRE](#)].
- [31] A. Celentano, L. Darmé, L. Marsicano and E. Nardi, *New production channels for light dark matter in hadronic showers*, *Phys. Rev. D* **102** (2020) 075026 [[arXiv:2006.09419](#)] [[INSPIRE](#)].
- [32] A. Berlin, S. Gori, P. Schuster and N. Toro, *Dark Sectors at the Fermilab SeaQuest Experiment*, *Phys. Rev. D* **98** (2018) 035011 [[arXiv:1804.00661](#)] [[INSPIRE](#)].
- [33] SHiP collaboration, *Sensitivity of the SHiP experiment to light dark matter*, *JHEP* **04** (2021) 199 [[arXiv:2010.11057](#)] [[INSPIRE](#)].
- [34] S. Gardner, R.J. Holt and A.S. Tadepalli, *New Prospects in Fixed Target Searches for Dark Forces with the SeaQuest Experiment at Fermilab*, *Phys. Rev. D* **93** (2016) 115015 [[arXiv:1509.00050](#)] [[INSPIRE](#)].

- [35] Y. Kahn, G. Krnjaic, J. Thaler and M. Toups, *DAEδALUS and dark matter detection*, *Phys. Rev. D* **91** (2015) 055006 [[arXiv:1411.1055](#)] [[INSPIRE](#)].
- [36] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [[INSPIRE](#)].
- [37] J.M. Berryman, A. de Gouvêa, P.J. Fox, B.J. Kayser, K.J. Kelly and J.L. Raaf, *Searches for Decays of New Particles in the DUNE Multi-Purpose Near Detector*, *JHEP* **02** (2020) 174 [[arXiv:1912.07622](#)] [[INSPIRE](#)].
- [38] GEANT4 collaboration, *GEANT4—a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250 [[INSPIRE](#)].
- [39] L. Buoncore, C. Frugiuele, F. Maltoni, O. Mattelaer and F. Tramontano, *Event generation for beam dump experiments*, *JHEP* **05** (2019) 028 [[arXiv:1812.06771](#)] [[INSPIRE](#)].
- [40] C. Degrande, C. Duhr, B. Fuks, D. Grellscheid, O. Mattelaer and T. Reiter, *UFO - The Universal FeynRules Output*, *Comput. Phys. Commun.* **183** (2012) 1201 [[arXiv:1108.2040](#)] [[INSPIRE](#)].
- [41] E. Fermi, *On the Theory of the impact between atoms and electrically charged particles*, *Z. Phys.* **29** (1924) 315 [[INSPIRE](#)].
- [42] E.J. Williams, *Nature of the high-energy particles of penetrating radiation and status of ionization and radiation formulae*, *Phys. Rev.* **45** (1934) 729 [[INSPIRE](#)].
- [43] C.F. von Weizsacker, *Radiation emitted in collisions of very fast electrons*, *Z. Phys.* **88** (1934) 612 [[INSPIRE](#)].
- [44] J. Blümlein and J. Brunner, *New Exclusion Limits on Dark Gauge Forces from Proton Bremsstrahlung in Beam-Dump Data*, *Phys. Lett. B* **731** (2014) 320 [[arXiv:1311.3870](#)] [[INSPIRE](#)].
- [45] P. deNiverville, C.-Y. Chen, M. Pospelov and A. Ritz, *Light dark matter in neutrino beams: production modelling and scattering signatures at MiniBooNE, T2K and SHiP*, *Phys. Rev. D* **95** (2017) 035006 [[arXiv:1609.01770](#)] [[INSPIRE](#)].
- [46] A. Faessler, M.I. Krivoruchenko and B.V. Martemyanov, *Once more on electromagnetic form factors of nucleons in extended vector meson dominance model*, *Phys. Rev. C* **82** (2010) 038201 [[arXiv:0910.5589](#)] [[INSPIRE](#)].
- [47] D.E. Morrissey and A.P. Spray, *New Limits on Light Hidden Sectors from Fixed-Target Experiments*, *JHEP* **06** (2014) 083 [[arXiv:1402.4817](#)] [[INSPIRE](#)].
- [48] D. Gorbunov, A. Makarov and I. Timiryasov, *Decaying light particles in the SHiP experiment: Signal rate estimates for hidden photons*, *Phys. Rev. D* **91** (2015) 035027 [[arXiv:1411.4007](#)] [[INSPIRE](#)].
- [49] C. Andreopoulos et al., *The GENIE Neutrino Monte Carlo Generator*, *Nucl. Instrum. Meth. A* **614** (2010) 87 [[arXiv:0905.2517](#)] [[INSPIRE](#)].
- [50] DUNE collaboration, *Simulated dune neutrino fluxes*.
- [51] DUNE collaboration, *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF*, (2015) [[arXiv:1512.06148](#)] [[INSPIRE](#)].
- [52] BABAR collaboration, *Search for Invisible Decays of a Dark Photon Produced in e^+e^- Collisions at BaBar*, *Phys. Rev. Lett.* **119** (2017) 131804 [[arXiv:1702.03327](#)] [[INSPIRE](#)].

- [53] D. Banerjee et al., *Dark matter search in missing energy events with NA64*, *Phys. Rev. Lett.* **123** (2019) 121801 [[arXiv:1906.00176](#)] [[INSPIRE](#)].
- [54] J. Beacham et al., *Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report*, *J. Phys. G* **47** (2020) 010501 [[arXiv:1901.09966](#)] [[INSPIRE](#)].
- [55] B. Batell, R. Essig and Z. Surujon, *Strong Constraints on Sub-GeV Dark Sectors from SLAC Beam Dump E137*, *Phys. Rev. Lett.* **113** (2014) 171802 [[arXiv:1406.2698](#)] [[INSPIRE](#)].
- [56] P. Fileviez Perez and M.B. Wise, *Breaking Local Baryon and Lepton Number at the TeV Scale*, *JHEP* **08** (2011) 068 [[arXiv:1106.0343](#)] [[INSPIRE](#)].
- [57] B.A. Dobrescu and C. Frugiuele, *Hidden GeV-scale interactions of quarks*, *Phys. Rev. Lett.* **113** (2014) 061801 [[arXiv:1404.3947](#)] [[INSPIRE](#)].
- [58] J.A. Dror, R. Lasenby and M. Pospelov, *Dark forces coupled to nonconserved currents*, *Phys. Rev. D* **96** (2017) 075036 [[arXiv:1707.01503](#)] [[INSPIRE](#)].
- [59] L. Michaels and F. Yu, *Probing new U(1) gauge symmetries via exotic $Z \rightarrow Z'\gamma$ decays*, *JHEP* **03** (2021) 120 [[arXiv:2010.00021](#)] [[INSPIRE](#)].
- [60] NA61/SHINE collaboration, *Measurements of Cross Sections and Charged Pion Spectra in Proton-Carbon Interactions at 31 GeV/c*, *Phys. Rev. C* **84** (2011) 034604 [[arXiv:1102.0983](#)] [[INSPIRE](#)].
- [61] PARTICLE DATA GROUP collaboration, *Review of Particle Physics*, *Phys. Rev. D* **98** (2018) 030001 [[INSPIRE](#)].
- [62] R.D. Ball et al., *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244 [[arXiv:1207.1303](#)] [[INSPIRE](#)].
- [63] NNPDF collaboration, *Parton distributions with QED corrections*, *Nucl. Phys. B* **877** (2013) 290 [[arXiv:1308.0598](#)] [[INSPIRE](#)].
- [64] MINIBOONE and MINOS collaborations, *First Measurement of ν_μ and ν_e Events in an Off-Axis Horn-Focused Neutrino Beam*, *Phys. Rev. Lett.* **102** (2009) 211801 [[arXiv:0809.2447](#)] [[INSPIRE](#)].
- [65] M.L. Graesser, I.M. Shoemaker and L. Vecchi, *A Dark Force for Baryons*, [arXiv:1107.2666](#) [[INSPIRE](#)].
- [66] BABAR collaboration, *A Search for Invisible Decays of the $\Upsilon(1S)$* , *Phys. Rev. Lett.* **103** (2009) 251801 [[arXiv:0908.2840](#)] [[INSPIRE](#)].
- [67] A. Atre, T. Han, S. Pascoli and B. Zhang, *The Search for Heavy Majorana Neutrinos*, *JHEP* **05** (2009) 030 [[arXiv:0901.3589](#)] [[INSPIRE](#)].
- [68] M. Drewes and B. Garbrecht, *Combining experimental and cosmological constraints on heavy neutrinos*, *Nucl. Phys. B* **921** (2017) 250 [[arXiv:1502.00477](#)] [[INSPIRE](#)].
- [69] A. de Gouvêa and A. Kobach, *Global Constraints on a Heavy Neutrino*, *Phys. Rev. D* **93** (2016) 033005 [[arXiv:1511.00683](#)] [[INSPIRE](#)].
- [70] P.D. Bolton, F.F. Deppisch and P.S. Bhupal Dev, *Neutrinoless double beta decay versus other probes of heavy sterile neutrinos*, *JHEP* **03** (2020) 170 [[arXiv:1912.03058](#)] [[INSPIRE](#)].
- [71] T. Boschi, *Nushock: Bunch of tools to study sensitivity of heavy neutrinos*, (2019) <https://github.com/tboschi/NuShock>.

- [72] C. Lourenco and H.K. Wohri, *Heavy flavour hadro-production from fixed-target to collider energies*, *Phys. Rept.* **433** (2006) 127 [[hep-ph/0609101](#)] [[INSPIRE](#)].
- [73] P.V. Ramana Murthy, C.A. Ayre, H.R. Gustafson, L.W. Jones and M.J. Longo, *Neutron Total Cross-Sections on Nuclei at Fermilab Energies*, *Nucl. Phys. B* **92** (1975) 269 [[INSPIRE](#)].
- [74] ZEUS collaboration, *Measurement of charm fragmentation fractions in photoproduction at HERA*, *JHEP* **09** (2013) 058 [[arXiv:1306.4862](#)] [[INSPIRE](#)].
- [75] S. Aoki et al., *Study of tau-neutrino production at the CERN SPS*, [arXiv:1708.08700](#) [[INSPIRE](#)].
- [76] Dawson-Haggerty et al., *trimesh* <https://trimsh.org/index.html>.
- [77] P. Ballett, M. Hostert, S. Pascoli, Y.F. Perez-Gonzalez, Z. Tabrizi and R. Zukanovich Funchal, *Neutrino Trident Scattering at Near Detectors*, *JHEP* **01** (2019) 119 [[arXiv:1807.10973](#)] [[INSPIRE](#)].
- [78] P. Ballett, S. Pascoli and M. Ross-Lonergan, *MeV-scale sterile neutrino decays at the Fermilab Short-Baseline Neutrino program*, *JHEP* **04** (2017) 102 [[arXiv:1610.08512](#)] [[INSPIRE](#)].
- [79] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) 1554 [*Erratum ibid.* **73** (2013) 2501] [[arXiv:1007.1727](#)] [[INSPIRE](#)].