

Shedding light on the MiniBoone Excess with Searches at the LHC

Christian Herwig,¹ Joshua Isaacson,¹ Bo Jayatilaka,¹ and Pedro A. N. Machado¹

¹*Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*

(Dated: September 12, 2023)

The origin of the excess of low-energy events observed by the MiniBooNE experiment remains a mystery, despite exhaustive investigations of backgrounds and a series of null measurements from complementary experiments. One intriguing explanation is the production of beyond-the-Standard-Model particles that could mimic the experimental signature of additional ν_e appearance seen in MiniBooNE. In one proposed mechanism, muon neutrinos up-scatter to produce a new “dark neutrino” state that decays by emitting highly-collimated electron-positron pairs. We propose high-energy neutrinos produced from W boson decays at the Large Hadron Collider as an ideal laboratory to study such models. Simple searches for a low-mass, boosted di-lepton resonance produced in association with a high- p_T muon from the W decay with Run 2 data would already provide unique sensitivity to a range of interesting dark neutrino scenarios. Looking farther ahead, we show how the unprecedented sample of W boson decays anticipated at the HL-LHC, together with improved lepton acceptance, and more sophisticated analysis techniques would explore much of the parameter space most compatible with the MiniBooNE excess.

Contents

I. Introduction	1
II. Dark Neutrino Model	3
III. LHC Search strategy	4
IV. Simulated samples and detector-level reconstruction	5
V. Analysis of dark neutrino production in W boson decays	7
A. Event selection	7
B. Signal extraction procedure and uncertainties	9
C. Analysis of displaced Z_D decays	11
VI. Results	12
VII. Conclusions and Outlook	13
References	14

I. INTRODUCTION

For several decades, the MiniBooNE experiment has presented an unresolved anomaly in its data, characterized by the observation of an excess of electron-like events over the predicted background in a muon-neutrino-dominated beam [1, 2]. Motivated by other short baseline neutrino anomalies, such as the LSND result [3], or the gallium anomaly [4], several proposed explanations for the MiniBooNE low energy excess

rely on oscillations or neutrino flavor transitions in general. Nevertheless, these explanations typically face strong constraints from other neutrino oscillation experiments [5–7], or from cosmological observations [8].

Due to these difficulties, other directions on how to address the MiniBooNE anomaly have been explored [9–16]. These use the fact that what MiniBooNE actually measures is the Cherenkov light produced by charged particles as they traverse its mineral oil detector. Therefore, signals that could produce electron-like Cherenkov rings have the potential to explain the anomaly.

Despite many attempts to explain the anomaly, its resolution is still an open problem. One intriguing model that has recently been proposed as a possible solution is the “dark neutrino” model [10, 11, 13]. In this scenario, the breaking of a new gauge symmetry acting only in Standard Model (SM) singlets, i.e., in the dark sector, leads to the smallness of neutrino masses. This mechanism naturally leads to mixing between standard model neutrinos and dark neutrinos, which generates a small interaction among these particles and the new gauge boson. This new interaction leads to a novel phenomenology at MiniBooNE: muon neutrinos can up-scatter to dark neutrinos via this new interaction, followed by the decay of the dark neutrinos into a light neutrino and an electron-positron pair. If the e^+e^- pair is sufficiently collimated, or if there is a large energy asymmetry in these final state particles, the MiniBooNE detector would classify it as an electron-like signal, thus providing a possible explanation for the anomaly.

The MicroBooNE experiment released their first set of results attempting to further investigate the MiniBooNE anomaly [17–20]. In these analyses, MicroBooNE requires the final state to have exactly one electron/positron. For the dark neutrino model, while the e^+e^- pair would be reconstructed as an electron-like signal at MiniBooNE, the signature at MicroBooNE could be quite different, since liquid argon detectors have much better particle reconstruction capabilities. Therefore, while the results from MicroBooNE are a great first step in addressing the MiniBooNE anomaly, they do not rule out the sterile neutrino interpretation of the MiniBooNE excess [21, 22], neither models with multiple lepton final states, nor other possible explanations of the excess [23].

From a theoretical view, the dark neutrino model is an interesting model of low-scale neutrino mass that could lead to observable phenomenology in laboratory experiments [11, 24]. Given the lack of experimental guidance on the scale of the physics that is responsible for neutrino mass generation, it is crucial to develop searches that could be sensitive to both high-scale models, such as GUT-embedded seesaw scenarios, or low-scale models such as the dark neutrino scenario. Although the dark neutrino phenomenology has been studied in neutrino and meson decay experiments, possible signatures at the LHC have not received much attention.

Dark neutrinos could be copiously produced by the decay of on-shell W bosons. With a W production cross-section of $\sigma(pp \rightarrow W) = 190$ nb at $\sqrt{s} = 13$ TeV center of mass energy [25], the LHC has produced roughly 30 billion W bosons during Run 2 data-taking at each of the ATLAS and CMS experiments. The ultimate High Luminosity LHC (HL-LHC) should enlarge these samples by an additional factor of 20-30, collecting an integrated luminosity up to 4 ab^{-1} and profiting from the slightly larger cross section at $\sqrt{s} = 14$ TeV. Therefore, even a small active-dark neutrino mixings could lead to a significant production of these new particles at the LHC. The dark neutrino will then decay into an e^+e^- pair, which would be the smoking gun signature of this model. While probing the very low-mass end of this scenario may be challenging at the LHC due to a large off-shell photon-to- e^+e^- background the LHC provides a complementary probe of such mass models, being sensitive to the hundred-MeV to multi-GeV scale. **PM: In this mass range, other decay modes could also open up, such as dimuon or hadronic final states.**

In this work, we explore the connection between LHC physics and neutrino experiments, highlighting an interesting complementarity between these high and low energy scale probes of new physics. We present a detailed study of the dark neutrino model in the context of LHC searches, showing how the high luminosity LHC can probe models of light new physics. We develop a new LHC search that can probe parameter space relevant to the MiniBooNE excess, and set leading constraints for a significant portion.

The paper is structured as follows Sec. II reviews the dark neutrino model and Sec. III details the search strategy at the LHC. The simulation of events along with event reconstruction is discussed in Sec. IV. An

Field	$SU(2)_L$	$U(1)_Y$	$U(1)_D$
N_R	1	0	+1
N_L	1	0	-1
H_1	2	0	+1
ϕ_1	1	0	+1
ϕ_2	1	0	+2

TABLE I: Field content of the dark neutrino scenario. The standard model is not charged under the new symmetry $U(1)_D$.

analysis of the possible constraints that the LHC could set on the dark neutrino model is carried out in Sec. V. Finally, Sec. VII provides some outlook to other possible searches that could be carried out in this low energy region and summarizes the findings of this work.

II. DARK NEUTRINO MODEL

The dark neutrino model consists of a realization of the neutrino mass mechanism at low scales [11]. In a nutshell, a new $U(1)_D$ gauge symmetry acting only on standard model singlets would lead to the conservation of lepton number if left unbroken. These “dark neutrinos” are vector-like under this new symmetry to ensure anomaly cancellation. The spontaneous breaking of this symmetry by scalar fields in the dark sector simultaneously breaks lepton number, which leads to nonzero neutrino masses. The breaking of the dark symmetry mixes active and dark neutrinos, leading to a new interaction of active neutrinos: an up-scattering to dark neutrinos via the exchange of the new gauge boson, Z_D , which kinetically mix with the photon. This new interaction could potentially explain the MiniBooNE anomaly [10]: the dark neutrinos would decay to light neutrinos and e^+e^- pairs; if the e^+e^- opening angle is small enough, MiniBooNE would classify these events as electron-like events, *i.e.* the signature of the anomaly.

To be more concrete, let us describe the Lagrangian of the dark neutrino scenario. The field content of the theory can be found in Table I. The mass Lagrangian is

$$\mathcal{L}_{\text{mass}} = y_\nu \bar{L} \tilde{H}_1 N_R + M \bar{N}_L N_R + y_D \phi_2^* \bar{N}_R^c N_R + y'_D \phi_2 \bar{N}_L^c N_L, \quad (1)$$

where we denote the left- and right-handed dark neutrino fields by N_L and N_R . After $U(1)_D$ breaking, this gives rise to an inverse seesaw texture

$$M_\nu = \begin{pmatrix} 0 & y_\nu \langle H_1 \rangle & 0 \\ y_\nu^\dagger \langle H_1 \rangle & y_D \langle \phi_2 \rangle & M \\ 0 & M & y'_D \langle \phi_2 \rangle \end{pmatrix}, \quad (2)$$

in the basis (ν_L, N_R^c, N_L) . Note that we have taken the vacuum expectation values (vevs) of the new scalar fields to be real. This leads to active neutrino masses of the order

$$m_\nu \sim y'_D \langle \phi_2 \rangle \frac{y_\nu^2 \langle H_1 \rangle^2}{M^2}, \quad (3)$$

PM: as long as $y'_D \langle \phi_2 \rangle \ll y_\nu \langle H_1 \rangle \ll M$, while the mixing between active neutrinos and N_R^c is given by $\theta \sim y_\nu^2 \langle H_1 \rangle^2 / M^2$, effectively decorrelating mixing from masses, as is usual in inverse seesaw scenarios. The ϕ_1 field is necessary to induce small vevs to both H_1 and ϕ_2 (see Ref. [11] for details).

From a phenomenological perspective, the mixing among flavor eigenstates gives rise to a coupling between the weak gauge bosons and the neutrino mass eigenstates, namely

$$\mathcal{L}_{\text{int}} = \frac{g}{\sqrt{2}} \sum_{\alpha} \sum_{i=1}^{e,\mu,\tau} U_{\alpha i} \bar{\ell}_{\alpha} \gamma_{\mu} P_L \nu_i W^{\mu} + \text{h.c.} + \frac{g}{c_W} \sum_{\alpha} \sum_{i=1}^{e,\mu,\tau} \sum_{j=1}^n U_{\alpha i}^* U_{\alpha j} \bar{\nu}_i \gamma_{\mu} P_L \bar{\nu}_j Z^{\mu}, \quad (4)$$

where g is the weak coupling constant, c_W is the cosine of the weak mixing angle, and $U_{\alpha i}$ is the $3 \times n$ PMNS matrix, where n is the total number of light and heavy neutrino states. For simplicity, we will assume that the PMNS matrix is mostly standard except for $U_{\mu 4}$ only. Denoting the heavy neutrino mass eigenstate as N_D , the branching ratio of the W boson to the new decay mode becomes

$$\mathcal{BR}(W \rightarrow \mu N_D) \simeq |U_{\mu 4}|^2 \mathcal{BR}(W \rightarrow \mu \nu), \quad (5)$$

assuming $|U_{\mu 4}|^2 \ll 1$ and $m_{N_D} \ll m_W$. The heavy neutrino then would decay $N_D \rightarrow Z_D \nu$ via the same mixing $U_{\mu 4}$, followed by $Z_D \rightarrow e^+ e^-$ due to the kinetic mixing with the photon ε . The width of N_D is given by, in the limit $m_{Z_D} \ll m_{N_D}$,

$$\Gamma(N_D \rightarrow \nu Z_D) \simeq \frac{g_D^2 |U_{\mu 4}|^2 m_{N_D}^3}{8\pi m_{Z_D}^2} = \frac{g_D^2}{45 \text{ nm}} \left(\frac{|U_{\mu 4}|^2}{10^{-7}} \right) \left(\frac{m_{N_D}}{100 \text{ MeV}} \right) \left(\frac{30 \text{ MeV}}{m_{Z_D}} \right)^2. \quad (6)$$

In this regime, for the parameters of interest, the N_D decay can be always considered prompt. For Z_D , on the other hand, the width depends on kinetic mixing with the photon, originating from the term

$$\mathcal{L}_{\text{km}} = \frac{\varepsilon}{2} F^{\mu\nu} F'_{\mu\nu}, \quad (7)$$

where F and F' are the photon and dark photon field strengths, and ε is the kinetic mixing parameter. The partial width to $e^+ e^-$ is, in the limit $m_e \ll m_{Z_D}$,

$$\Gamma(Z_D \rightarrow e^+ e^-) \simeq \frac{\alpha \varepsilon^2}{3} m_{Z_D} = \frac{1}{0.1 \text{ mm}} \left(\frac{\alpha \varepsilon^2}{2 \times 10^{-10}} \right) \left(\frac{m_{Z_D}}{30 \text{ MeV}} \right), \quad (8)$$

where α is the fine structure constant. For the parameters of interest, the decay of Z_D happens within the inner tracker in nearly all cases. While a heavy Z_D would generally decay promptly, lighter bosons would experience larger Lorentz boosts from the W decay leading to macroscopic displacements (e.g. 1 mm to 10 cm) and thus a displaced vertex signature. For a heavier Z_D , other decay channels will open up, and the partial widths can be properly obtained using the $R(s)$ ratio, that is, the ratio between the cross sections for $e^+ e^-$ to hadrons versus muons, as in Ref. [26]. We have implemented the dark neutrino model in FeynRules [27, 28] and created the set of UFO [29, 30] files needed for use within BSM event generator tools. The `DarkNeutrino` model is made publicly available on Zenodo [31].

III. LHC SEARCH STRATEGY

The dark neutrino model described above opens two potential portals to produce dark sector particles from their SM counterparts, through either mixing of the neutral leptons or $U(1)$ gauge bosons. This suggests the potential for multiple search strategies to contribute complementary sensitivity, including experiments that employ a range of particle species and energies. In addition to the high-intensity neutrino sources like the FNAL Booster Neutrino Beam that serviced the MiniBooNE experiment, beam-dump and fixed-target setups can similarly explore this model through neutrinos produced in light mesons and muon decays [32, 33]. Alternatively, the $U(1)_D$ portal can be explored through a range of experiments targeting dark photon production and prompt decay, including $e^+ e^-$ colliding beams, fixed targets, and meson decays. High-energy proton-proton collisions can also present powerful constraints on this possibility at high mass [34]. Specific constraints depend on the dominant portal coupling that may be realized in Nature, in addition to the masses of the BSM states. For the range of parameters considered in this work, the dark neutrino decays promptly, so that constraints on long-lived heavy neutral leptons do not apply.

The LHC provides a large dataset of neutrinos across a range of energies that may be exploited to test scenarios where the neutral lepton mixing parameter $U_{\mu 4}$ is significant **PM: what dowe mean by this?**. At




FIG. 1: Representative diagrams are shown for the production of e^+e^- pairs from dark photon decays in $W^- \rightarrow \mu^- \bar{\nu}_\mu$ events, produced either through the mixing of neutral leptons (left) or $U(1)$ gauge forces (center). At right the dominant SM background process is shown, in which the muon radiates an e^+e^- pair via virtual photon emission.

the ATLAS and CMS experiments, neutrinos from decays of the massive weak bosons provide a promising source, most notably the W^\pm . With a large cross section of 190 nb at 13 TeV proton-proton center of mass energy, nearly 30 billion W bosons were produced at each of the ATLAS and CMS interaction points over the course of Run 2 ($\sim 150 \text{ fb}^{-1}$) of the LHC [25, 35]. Leptonic decays $W^\pm \rightarrow \ell^\pm \nu$ ($\ell = e$ or μ) present a promising experimental channel signature because of the presence of an energetic charged lepton that can be used to trigger and cleanly reconstruct the event. Figure 1 (left) shows a diagram corresponding to signal production through this process, where a dark neutrino is produced and decays to two charged leptons and a SM neutrino. The same final state can access the dark photon portal through diagrams such as the one shown in Figure 1 (center), though it is generally subdominant for the combination of model parameters considered in this work. Figure 1 (right) shows a representative diagram for the dominant SM background process to this signature, where a virtual photon is radiated off of the charged lepton from the W decay.

In the following sections, the potential sensitivity of this channel is investigated using Monte Carlo simulation to compare the dark neutrino and SM processes and devise a promising search strategy. The key feature to be targeted is the low-mass pair of charged leptons produced through the dark photon decay. This resonance provides a clean experimental signature and allows a robust “bump hunt” strategy to extract the peaking signal over potential background processes though the use of sideband techniques. For light Z_D masses, the macroscopic displacement of the lepton pair can also be used as an effective method of background rejection.

Because the Z_D decays via coupling to SM hypercharge, the spectrum of possible final states depends only on its mass. The branching fraction to leptons generally increases for shrinking Z_D masses, with significant exceptions near the hadronic resonances that can be calculated using $R(s)$ ratio data [26]. As an example, $\mathcal{BR}(Z_D \rightarrow e^+e^-) \sim 15\%$ at 10 GeV, growing to nearly 50% at 400 MeV. Below this value, the fraction of decays to the electron and muon channels begin to diverge, with $\mathcal{BR}(Z_D \rightarrow e^+e^-) = 100\%$ below $2m_\mu$. Unless otherwise specified we take the sum of the electron and muon channels as the signal in the remainder of this work.

IV. SIMULATED SAMPLES AND DETECTOR-LEVEL RECONSTRUCTION

Simulated samples of Monte Carlo events corresponding to SM processes and the dark neutrino signals are used to study analysis selection criteria and establish sensitivity to the BSM process. SM backgrounds are calculated using MG5_AMC@NLO 3.4.1 at LO, interfaced to PYTHIA 8.306, including the dominant process of inclusive $\ell^\pm \bar{\nu}_\ell e^+e^-$ production. Leptons are generated with $p_T > 1 \text{ GeV}$ and $|\eta| < 5$, and requiring that the mass of the same-flavor, opposite-sign dilepton pair m_{ee} is greater than 10 MeV. In the inclusive phase space, the leading contribution to this process comes when a virtual photon is radiated off the charged lepton resulting from the W decay, followed by radiation off of an initial-state quark. The process

is normalized by applying a k -factor corresponding to the full NNLO + NLO electroweak charged-current prediction [35]. The corresponding process with a Z boson produced in association with a pair of low-mass electrons in place of a W is also simulated at LO and normalized to the NNLO calculation [36–38] in a similar manner. Lastly, a MC calculation corresponding to the production of top-quark pairs in association with a soft virtual $Z/\gamma \rightarrow \ell^+\ell^-$ is used, modified by the ratio of NNLO+NNLL to LO cross section for top-quark pairs [39, 40]. The expected contribution from rarer SM processes yielding three or more prompt leptons is negligible for the phase space considered in this analysis.

Signal samples are generated with the same tools, using the Dirac DARKNEUTRINO model introduced in Section II. **PM: Our results depend weakly on the Dirac vs. Majorana nature of the dark neutrino.** Dark neutrinos produced via W^\pm decay are simulated from the $\ell^-\bar{\nu}_\ell Z_D$ matrix element, allowing Pythia to decay the Z_D to leptons. As a baseline selection of parameters, we set **PM: changed to U^2 :** $|U_{\mu 4}|^2 = 10^{-4}$, $\alpha\epsilon^2 = 2 \times 10^{-10}$, and $g_D = 0.25$, and consider masses $m_{Z_D} < m_{N_D}$ corresponding to the prompt 2-body decay of the N_D . In this configuration $\ell^-\bar{\nu}_\ell Z_D$ production is dominated by the neutrino portal rather than direct radiation of a Z_D off an initial-state quark (though the latter can become important for very small values of $|U_{\mu 4}|^2$) and the cross section is largely insensitive to ϵ and g_D . **PM: I am confused about this statement.** To account for higher-order effects that are not included in the signal calculation, the k -factor for the SM W^\pm process is also applied to signal. At $\sqrt{s} = 13$ TeV the resulting cross section is 2.1 pb for when $m_{N_D} = 300$ MeV and $m_{Z_D} = 100$ MeV, but varies within less than 10% when $1 \text{ MeV} < m_{Z_D} < m_{N_D} < 10 \text{ GeV}$.

Simulated particle-level quantities are translated to detector-level observables with a simplified detector simulation in order to estimate the potential reach of the ATLAS and CMS experiments to the signatures above. Because the key characteristics of the signal are encoded in the kinematics of the Z_D , we restrict ourselves to the consideration of leptonic observables only in this work. While improved sensitivity may be achieved by considering other features of the collision events (such as the imbalance of the transverse momenta of all reconstructed particles), these can depend strongly on the details of the experimental resolution and running conditions, and are thus omitted from the present study.

High-energy dark photons produced through the signal processes described above lead to pairs of high-energy particles that must be exceptionally collimated in the case of a low-mass Z_D . Consequently, the resolution on the Z_D dilepton mass $m_{\ell\ell}$, which drives the background rate, relies on the precision of the leptons' angular reconstruction. We apply a p_T -dependent Gaussian smearing on reconstructed electron kinematics, corresponding to 0.001-0.0025 radians in ϕ and 0.0005-0.002 in units of $\cot\theta$. In addition, a 2.5-8% smearing on the electron p_T is assessed, taken as a conservative estimate in the phase space most relevant to this search[41, 42].¹ Figure 2 shows the mass resolution that is obtained for pairs of electrons with these smearing functions applied, as evaluated on signal MC samples with N_D mass fixed to 10 GeV. The variation found by considering the upper and lower extremes of the single-object kinematic resolutions correspond to the region enclosed in the $\pm 1\sigma$ band. In the muon decay channel, the mass resolution is well-established from measurements of light meson decays, which we conservatively take as 2%. A 15% inefficiency factor per lepton is also applied to account for the effect of requirements on the reconstruction quality entering through identification and isolation cuts. While most events with four leptons within the detector acceptance described above should be removable via loose lepton or track vetos, we assume a 5% inefficiency factor that allows these events to populate our signal region.

The lab-frame displacement of the Z_D is calculated for signal events, based on the proper lifetime (proportional to $1/\epsilon^2 m_{Z_D}$) and large momentum inherited from its production in the W boson decay. For the analysis considering SM backgrounds with prompt leptons, signal events are required to have displacements less than 1 mm. For signal events with displacements from 1 mm to 10 cm, the SM backgrounds described

¹ We adopt standard collider reference coordinate system centered at the interaction point with the z -axis oriented along the beamline, the x -axis pointing to the center of the LHC ring, and the y -axis pointing upward. The (r, ϕ) plane is transverse to the beam, with θ describing the polar angle from $+\hat{z}$.




FIG. 2: The di-electron mass resolution is shown for electrons for the optimistic and pessimistic scenarios described in the text. Single-electron resolutions are applied to simulated events in the dark neutrino signal model for a 10 GeV N_D mass. The quoted resolution is computed as the smallest interval containing 68% of the smeared mass values.

above can be efficiently rejected based on the separation of the dilepton vertex from the primary vertex, which can be well-reconstructed from the high- p_T W -decay lepton and the recoil system. In this case it is reasonable to expect that a background-free search can be conducted, at the cost of some additional penalty of the signal efficiency described in Section ??.

V. ANALYSIS OF DARK NEUTRINO PRODUCTION IN W BOSON DECAYS

To maximize the sensitivity to small neutrino portal couplings, we aim to construct a relatively inclusive analysis retaining high signal efficiency. Thus, typical W boson progenitors of the BSM signal will be produced nearly at rest, leading to a single high- p_T lepton in one hemisphere of the event opposite a low-mass, low-momentum di-lepton pair in the other. The dominant background process is W boson production involving the radiation of a virtual $Z/\gamma^* \rightarrow \ell^+\ell^-$. This can originate from ‘diboson-like’ processes, where the lepton pair is radiated from an initial-state quark or the W propagator, or final-state radiation off the charged lepton produced in the W decay. Before any selection is applied, the largest contribution comes from the latter class of processes, which can be effectively reduced by requiring the soft di-lepton pair to be well-separated from the third, high- p_T lepton.

Rarer backgrounds arise from analogous processes where a soft lepton pair is produced in association with a Z boson or a top and anti-top quark pair. A challenging source of background can arise from Z to four lepton events, when one of the $Z \rightarrow \ell^+\ell^-$ daughter leptons radiates another soft lepton pair if the high- p_T radiating daughter lepton fails to be reconstructed. While backgrounds with top quarks can lead to a similar final state, they can be suppressed by the accompaniment of hadronic activity. Other processes are either significantly rarer than those considered above or can be well-suppressed by similar requirements. Experimental backgrounds may also arise from the mis-identification of charged hadrons and non-prompt leptons and must be estimated directly from the data. Uncertainties stemming from this and other sources are discussed further in Section V B.

A. Event selection

A baseline set of events is defined for the analysis which should ensure the signal can be well-reconstructed by the experimental apparatus. Three leptons are required, each of which enters the central region of the detector ($|\eta| < 2.5$) and meet a minimum p_T requirement. These are kept as low as possible to maximize the signal rate from the light Z_D decay, taking 3 GeV muons and 5 GeV electrons as the




FIG. 3: The expected distribution of BSM signal and background processes are shown for events passing the pre-selection criteria, normalized to 150 fb^{-1} of collected data. SM background contributions are shown cumulatively for W , Z , and top processes, while realizations of the signal model for different Z_D and N_D masses are overlaid, fixing $|U_{\mu 4}|^2 = 10^{-4}$. Overflow values falling below (above) the x-axis limits are included in the first (last) bin.

baseline. A scenario including electrons with $p_T > 1 \text{ GeV}$ is also considered, motivated by more aggressive strategies being pursued by the CMS Experiment [43]. The majority of $W \rightarrow \ell\nu$ decay events can be collected through the use of triggers requiring a single high- p_T electron or muon at ATLAS and CMS. Thus, at least one of the three well-reconstructed leptons should also satisfy $p_T > 25 \text{ GeV}$, which is taken as a representative choice for the suite of trigger paths used across experiments and lepton flavors. The mass of the two leptons with the lowest transverse momenta $m_{\ell\ell}$ is required to be less than 10 GeV to focus on the region most consistent with the MiniBoone anomaly.

As the signal is a purely electroweak process, events with additional hadronic activity such as top quarks decays are vetoed. The sum of charged hadronic momenta H_T is defined as the p_T sum of all charged hadrons satisfying $p_T > 2 \text{ GeV}$ and $|\eta| < 2.5$ and is required to be less than 30 GeV .

Figure 3 shows a comparison of various signals and the cumulative SM background processes after these baseline requirements are enforced. The expected yields for W , Z , and top-quark processes are shown for 13 TeV proton-proton center of mass energy, normalized to 150 fb^{-1} of data, corresponding to the approximate sizes of the Run 2 data sets collected by each of the ATLAS and CMS Experiments from 2015-2018. The expected spectrum of events is also shown for the dark neutrino model for several choices of Z_D and N_D masses. Leptons are labeled l_1, l_2, l_3 in order of descending p_T .

Several key differences between the signal and BSM processes can be exploited to define a set of events enriched in potential dark neutrino decays. In the majority of SM W decay events, the soft di-lepton pair is most often radiated off of the charged lepton, as opposed to off of a neutrino in the BSM process. In addition to the large angular difference $d\phi(l_1, \ell^+ \ell^-)$ between the high- p_T and sub-leading leptons, this lowers the average momentum of the charged lepton (neutrino) from the initial W decay in background (signal) events. As a consequence, signal events pass the high- p_T lepton trigger more efficiently and lead

Observable	Run 2 selection
ℓ multiplicity	exactly three
Electron p_T	$> 1, 5$ GeV
Muon p_T	> 3 GeV
$ \eta_\ell $	< 2.5
$p_T(\ell_1)$	> 25 GeV
$d\phi(\ell_1, Z_D)$	> 2.7
$dR(\ell_1, Z_D)$	$\in [2.7, 3.5]$
$m(\ell_1, \ell_2, \ell_3)$	$< m_W$
$p_T(\ell_1, \ell_2, \ell_3)$	$< m_W/2$

TABLE II: Summary of the selection criteria, including baseline cuts in addition to selections optimized for the LHC Run 2 data sets. Scenarios with a minimum electron p_T of both 1 and 5 GeV are considered.

to a smaller imbalance in transverse momentum among the reconstructed visible objects ($p_{T,\text{miss}}$). The last effect can also clearly be seen in the magnitude of the vector sum of the three lepton momenta $p_T(\ell_1, \ell_2, \ell_3)$, equivalent to $p_{T,\text{miss}}$ in the limit where the W boson is produced at rest. The mass spectra of the the soft di-lepton pair is smoothly-falling for the SM background.

Table II outlines a set of selection criteria based on these differences, designed to significantly remove the background while retaining a high signal efficiency for all combinations of N_D and Z_D masses. The azimuthal difference between the leading lepton and the Z_D candidate formed by the remaining two leptons must be larger than 2.7 radians. Additionally, the distance $dR = \sqrt{d\phi^2 + d\eta^2}$ between the leading lepton and Z_D candidate is required to be between 2.7 and 3.5. The mass of the three leptons is required to be less than 80.3 GeV to be consistent with the W decay hypothesis. The p_T of the three-lepton system, equivalent to $p_{T,\text{miss}}$ in the limit of a W produced at rest should be less than half the W mass.

Figure 4 shows several properties of the remaining events that pass this selection. **PM: Note that we have adjusted the mixing to $|U_{\mu 4}|^2 = 10^{-6}$ for visual clarity.** At this point several key handles remain to discriminate signal from the SM background. While the mass of the di-lepton system is smoothly falling for the background, the signal process peaks at the Z_D mass. The value of the N_D mass also significantly impacts the signal kinematics. When the N_D is significantly heavier than the Z_D the lepton p_T spectra tends to peak at low values. When the mass difference between the new states becomes small, the majority of the N_D momenta is transferred to the leptons through the decay $N_D \rightarrow \nu(Z_D \rightarrow \ell^+\ell^-)$. As a consequence, less momentum is carried by the neutrino, the charged lepton p_T favor larger values, and the three-lepton mass peaks closer to m_W . While our projected limits are based on the reconstructed Z_D candidate mass alone, future work may exploit this information using more sophisticated techniques such as multivariate discriminants, kinematic fits, or matrix element methods to take advantage of these features.

B. Signal extraction procedure and uncertainties

The mass of the lepton pair from the Z_D decay has not been explicitly included in the selection described above, and can be used to extract signals in the presence of the smoothly falling SM background distribution. This feature allows the non-resonant background to be estimated from a fit using the ‘data side-band’ technique, whereby interpolation to $m_{\ell\ell}$ values populated by the signal, can commonly constrain the expected background to the sub-% level. A significance metric is calculated from the number of expected signal N_S and background events N_B in a di-lepton mass window as $S = N_S/\sqrt{N_B}$, where models that predict N_S corresponding to $S > S_{\text{excl}} \equiv 2$ are expected to be excluded. Because the mass resolution varies with both the Z_D candidate mass itself as well as the lepton flavor, as described in Section IV, an optimal window in $m_{\ell\ell}$ is used to extract N_S for each scenario under consideration. This significance-maximizing




FIG. 4: The expected distribution of BSM signal and background processes are shown for events passing the signal region selection, normalized to 150 fb^{-1} of collected data. SM background contributions are shown cumulatively for W , Z , and top processes, while realizations of the signal model for different Z_D and N_D masses are overlaid, fixing $|U_{\mu 4}|^2 = 10^{-6}$. Overflow values falling below (above) the x-axis limits are included in the first (last) bin.

interval ranges from 10 to 500 MeV for the electron channel and 3 to 100 MeV for the muon channel. In all scenarios considered, no $m_{\ell\ell}$ window is selected which contains fewer than 10 background events.

Because the production of N_D in the analysis phase space is dominated by W decays, the expected signal cross section is directly proportional to $|U_{\mu 4}|^2$. This allows a limit on the number of signal events to be translated to a limit on $U_{\mu 4}$ based on the significance $S_{\mu 4, \text{ref}}$ obtained with the sample of Monte Carlo events generated with a reference coupling $|U_{\mu 4, \text{ref}}|^2$ via the relation $|U_{\mu 4, \text{excl}}|^2 = (S_{\text{excl}}/S_{\text{ref}}) \cdot |U_{\mu 4, \text{ref}}|^2$. This extrapolation is verified by additional MC calculations of the cross sections and kinematic distributions with re-scaled coupling parameters.

This statistical procedure is repeated for modifications of the analysis accounting for the impact of various systematic uncertainties. Without a complete detector simulation it is difficult to assess the importance of backgrounds due to mis-identified and non-prompt leptons, including B decays and photon conversions. If lepton identification criteria cannot reduce these backgrounds to a negligible level, they may constitute an additional background that varies smoothly in $m_{\ell\ell}$, up to known hadronic resonances that can be masked. We address this possibility with a conservative, ad-hoc approach of scaling the background mass template by a factor of 1.5 and applying the full size of the correction as an uncertainty. The mass resolution also has a large impact on the result, as it effectively determines the number of background events overlapping with a given signal peak. For muons this is taken to be $2 \pm 1\%$ independent of mass whereas for electrons the smearing functions and their uncertainties are those described in Section IV. Uncertainties are also considered to account for the efficiency to select and identify leptons, taken to be $85 \pm 5\%$.




FIG. 5: At left, the distribution of displacements is shown for various combinations of signal model parameters, normalized to unity. The subset of events passing the full event selection is also shown for comparison. At right, the same distributions are shown for events passing the full event selection with variations of the kinetic mixing parameter ϵ . Events falling below or above of the displayed range of displacements are added to the first or last bin of displacement, respectively.

C. Analysis of displaced Z_D decays

In the case where the Z_D is light and weakly-enough coupled to become significantly long-lived in the lab frame, a parallel analysis strategy can be pursued to the prompt search described above. In order to select events that are well-reconstructed and compatible with the Dark Neutrino signal we consider the same selection requirements developed for the prompt analysis, with the additional requirement that all leptons have $p_T > 5$ GeV. Instead of requiring the pair of lower- p_T leptons to originate from a vertex with displacement less than 1 mm, distances of 1 mm to 10 cm are considered. The distribution of expected Z_D displacements is shown in Figure 5 for a range of representative signal model parameter sets. In the case of a high-mass Z_D , nearly all events populate the prompt analysis bin, whereas a wider distribution of lifetimes (extending beyond a meter in some cases) are found for the 30 MeV benchmark. While the selection requirements are found not to significantly affect the shape of the expected displacement, varying values of the kinetic mixing ϵ have a significant impact.

For displacements in the range of 1 mm to 10 cm, the lepton reconstruction efficiency should be excellent, with lepton tracks still expected to leave hits in all or nearly all of the silicon tracker layers. There is no SM process that leads to the expected signature of a pair of resonant leptons with a significant displaced vertex. Displaced, non-resonant lepton pairs from the decay of B mesons produced in the W recoil may be important for shorter lifetimes. Unfortunately the estimation of such fake/non-prompt backgrounds are outside the scope of this study, only being reliably estimated using data-driven methods by the LHC experiments. However analyses studying related signatures have demonstrated that these backgrounds can be effectively controlled at the required level in Run 2 (see, e.g. Refs. [1]), even before the requirement of a narrow di-lepton mass window. To estimate the approximate potential of the displaced analysis search strategy, we extract expected limits based on the requirement that at least fifty signal events pass the kinematic requirements. This value is chosen to allow an additional signal inefficiency of 10% in order to select a high-quality displaced di-lepton vertex, while still leaving a enough expected events to set a limit under the background-free assumption. Together with the signal efficiency of the kinematic requirements, this total efficiency goal is in line with that reported by similar analyses by ATLAS and CMS [1]. **Todo: add cites for this 0.1-1.5% signal efficiency number.**

CH: A few more last things to do here:

1. include full width (not partial ee) in the calculation




FIG. 6: The ranges of neutrino mixing parameters expected to be excluded by LHC experiments during Run 2 and the HL-LHC are shown versus N_D mass for a fixed Z_D mass of 30 MeV. Upper limits on $U_{\mu 4}$ are shown for both the prompt and displaced search strategies, with uncertainty bands described in the text. Closed contours correspond to the parameters favored by Ref. [10] at the 1, 2, 3, 4 and 5-sigma confidence level.

2. swap up/down epsilon values

VI. RESULTS

Expected limits are placed on the neutrino portal coupling $|U_{\mu 4}|^2$ for a range of signal mass hypotheses and experimental conditions. Limits are shown for the $\sqrt{s} = 13$ TeV center of mass energy and 150 fb^{-1} , corresponding to the potential reach of either ATLAS or CMS might place given data already recorded. These results are also extrapolated to a potential high-luminosity LHC scenario where 4 ab^{-1} are collected at a $\sqrt{s} = 14$ TeV center of mass energy. Because this difference in energy is relatively minor, the same sets of MC events are used, with their respective cross-sections scaled to their values at $\sqrt{s} = 14$ TeV.

Figure 6 shows projected exclusions for the case of a Z_D mass fixed to 30 MeV, varying m_{N_D} . Existing constraints are also displayed [32, 33] in addition to the region of parameters favored by Ref. [10]. Results are shown for both the prompt and displaced analyses, projected for both the Run 2 data set and the HL-LHC. Shaded bands indicate variations in the limit corresponding to uncertainties for the prompt search and variations in the signal efficiency for the displaced search, described in Sections V B and V C, respectively. Because of the Z_D is very light in this scenario and thus considerably boosted the displaced analysis is dominant, with sensitivity of the prompt analysis being similar to existing constraints.

Figure 7 shows excluded regions of N_D and $U_{\mu 4}$ under the assumption of $m_{N_D} = 3m_{Z_D}$. Expected limits are shown for both the prompt and displaced analysis strategies, separately for integrated luminosities and energies corresponding to Run 2 and the HL-LHC. The proper branching fractions of the Z_D are considered for the electron and muon channels, which are also statistically combined. For low Z_D masses, larger lifetimes and boost factors cause the displaced analysis to contribute the dominant sensitivity while at larger masses, the prompt analysis yields the strongest limit.

Differences in the value of kinematic mixing parameter ϵ can give rise to different expected distributions




FIG. 7: Expected upper limit on the neutrino mixing parameter $U_{\mu 4}$ are shown as a function of the Z_D mass under the assumption that $m_{N_D} = 3m_{Z_D}$. The results of the prompt and displaced analysis are separately presented, assuming both the Run 2 and HL-LHC data sets.

of the displaced vertex displacement for a fixed Z_D mass. These impact of these differences are shown in Figure 8, where the limits expected for the nominal value of $\alpha\epsilon^2 = \times 10^{-10}$ are compared to those for values of $\alpha\epsilon^2 = \times 10^{-9}$ and $\alpha\epsilon^2 = \times 10^{-11}$. **CH: confirm the final values.** For the scenario with a 30 MeV Z_D mass the displaced analysis continues to dominate in all cases, with the expected constraint varying based on the expected fraction of signal falling into the 1 mm to 10 cm search window. For the scenario with $m_{N_D} = 3m_{Z_D}$, an upward variation of ϵ generally will strengthen the limit from the prompt search and weaken the limit from the displaced one and vice-versa. However, for very light Z_D the nominal choice of ϵ leads to a signal peaking outside of the acceptance considered for the displaced analysis, so consequently larger values of the parameter will improve the limit for both the prompt and displaced search. **CH: Terminate the limits at the excluded A' lower masses.**

The expected lepton p_T spectra depends on the specific signal mass parameters under consideration, but generally peaks at low values, motivating the use of minimal thresholds for analysis. Figure 9 compares various possibilities along these lines. In the prompt analysis, the nominal scenario of 3 GeV muons and 5 GeV electrons is augmented by expected limits that would result from lowering the electron p_T threshold to 1 GeV. The displaced analysis nominally considers electrons and muons with p_T above 5 GeV, with a 1 GeV scenario also considered for electrons. In all cases, limits for the combined electron and muon channels are computed by adding the expected significance for each channel in quadrature.

VII. CONCLUSIONS AND OUTLOOK

Despite many significant efforts to understand the source of the excess of low-energy events seen by the MiniBooNE, a satisfying solution remains elusive. The dark neutrino model offers an exciting possibility that this anomaly could be our first hint of a rich new sector of particles, sterile under the SM gauge forces. **PM: , but with its own gauge symmetry and scalar sector.** We have proposed a new method to probe this compelling scenario using high-energy neutrinos from the large samples of W bosons collected




FIG. 8: Expected upper limit on the neutrino mixing parameter $U_{\mu 4}$ are shown as a function of the Z_D mass for several assumptions on the kinetic mixing parameter. Comparisons are made for the Run 2 (top row) and HL-LHC data sets (bottom row), and for the $m_{Z_D} = 30$ MeV (left column) and $m_{N_D} = 3m_{Z_D}$ (right column). In each case, the results of both the prompt and displaced searches are shown, assuming values of $\alpha\epsilon^2 = \times 10^{-9}$, $\alpha\epsilon^2 = \times 10^{-10}$, and $\alpha\epsilon^2 = \times 10^{-11}$.

by the ATLAS and CMS experiments at the LHC. Search strategies considering both prompt and displaced lepton pairs should each be pursued and contribute complementary sensitivities that depend on the particular parameters of the signal model. The coverage of these strategies is shown to be highly complementary with low-energy experiments, reaching mixing parameters $|U_{\mu 4}|^2$ from 10^{-6} to 10^{-8} across a broad range of dark neutrino masses. Notably, very low Z_D masses, accessible only via the displaced decays to electron pairs, can be explored thanks to the large Lorentz boost factor afforded by the W boson production channel. Unique sensitivity can be achieved with the $\sqrt{s} = 13$ TeV data that has already been collected, and the High-Luminosity LHC will provide an unprecedented sample of W decays to probe dark neutrino production with exceeding-small portal couplings.

Acknowledgments: We warmly thank Anadi Canepa for making the initial connections that lead to this work. We also thank Stefan H"och and Frank Krauss for useful conversations. The authors' work is supported by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

-
- [1] A. A. Aguilar-Arevalo et al. (MiniBooNE), Phys. Rev. Lett. **98**, 231801 (2007), 0704.1500.
 [2] A. A. Aguilar-Arevalo et al. (MiniBooNE), Phys. Rev. D **103**, 052002 (2021), 2006.16883.




FIG. 9: Expected upper limit on the neutrino mixing parameter $U_{\mu 4}$ are shown as a function of the Z_D mass for several assumptions about the classes of potentially reconstructable leptons. Comparisons are made for the Run 2 (top row) and HL-LHC data sets (bottom row), and for the $m_{Z_D} = 30$ MeV (left column) and $m_{N_D} = 3m_{Z_D}$ (right column).

- [3] A. Aguilar et al. (LSND), Phys. Rev. D **64**, 112007 (2001), hep-ex/0104049.
[4] V. V. Barinov et al., Phys. Rev. C **105**, 065502 (2022), 2201.07364.
[5] G. H. Collin, C. A. Argüelles, J. M. Conrad, and M. H. Shaevitz, Nucl. Phys. B **908**, 354 (2016), 1602.00671.
[6] S. Gariazzo, C. Giunti, M. Laveder, and Y. F. Li, JHEP **06**, 135 (2017), 1703.00860.
[7] M. Dentler, A. Hernández-Cabezudo, J. Kopp, P. A. N. Machado, M. Maltoni, I. Martinez-Soler, and T. Schwetz, JHEP **08**, 010 (2018), 1803.10661.
[8] N. Aghanim et al. (Planck), Astron. Astrophys. **641**, A6 (2020), [Erratum: Astron.Astrophys. 652, C4 (2021)], 1807.06209.
[9] S. N. Gninenko, Phys. Rev. Lett. **103**, 241802 (2009), 0902.3802.
[10] E. Bertuzzo, S. Jana, P. A. N. Machado, and R. Zukanovich Funchal, Phys. Rev. Lett. **121**, 241801 (2018), 1807.09877.
[11] E. Bertuzzo, S. Jana, P. A. N. Machado, and R. Zukanovich Funchal, Phys. Lett. B **791**, 210 (2019), 1808.02500.
[12] C. A. Argüelles, M. Hostert, and Y.-D. Tsai, Phys. Rev. Lett. **123**, 261801 (2019), 1812.08768.
[13] P. Ballett, S. Pascoli, and M. Ross-Lonergan, Phys. Rev. D **99**, 071701 (2019), 1808.02915.
[14] O. Fischer, A. Hernández-Cabezudo, and T. Schwetz, Phys. Rev. D **101**, 075045 (2020), 1909.09561.
[15] W. Abdallah, R. Gandhi, and S. Roy, JHEP **12**, 188 (2020), 2006.01948.
[16] B. Dutta, D. Kim, A. Thompson, R. T. Thornton, and R. G. Van de Water, Phys. Rev. Lett. **129**, 111803 (2022), 2110.11944.
[17] P. Abratenko et al. (MicroBooNE), Phys. Rev. Lett. **128**, 241801 (2022), 2110.14054.
[18] P. Abratenko et al. (MicroBooNE), Phys. Rev. D **105**, 112005 (2022), 2110.13978.
[19] P. Abratenko et al. (MicroBooNE), Phys. Rev. D **105**, 112003 (2022), 2110.14080.
[20] P. Abratenko et al. ((The MicroBooNE Collaboration)*, MicroBooNE), Phys. Rev. D **105**, 112004 (2022), 2110.14065.

- [21] C. A. Argüelles, I. Esteban, M. Hostert, K. J. Kelly, J. Kopp, P. A. N. Machado, I. Martinez-Soler, and Y. F. Perez-Gonzalez, *Phys. Rev. Lett.* **128**, 241802 (2022), 2111.10359.
- [22] P. Abratenko et al. (MicroBooNE), *Phys. Rev. Lett.* **130**, 011801 (2023), 2210.10216.
- [23] C. A. Argüelles et al., *Rept. Prog. Phys.* **83**, 124201 (2020), 1907.08311.
- [24] P. Ballett, M. Hostert, and S. Pascoli, *Phys. Rev. D* **101**, 115025 (2020), 1903.07589.
- [25] G. Aad et al. (ATLAS), *Phys. Lett. B* **759**, 601 (2016), 1603.09222.
- [26] D. Curtin, R. Essig, S. Gori, and J. Shelton, *JHEP* **02**, 157 (2015), 1412.0018.
- [27] N. D. Christensen and C. Duhr, *Comput. Phys. Commun.* **180**, 1614 (2009), 0806.4194.
- [28] N. D. Christensen, P. de Aquino, C. Degrande, C. Duhr, B. Fuks, M. Herquet, F. Maltoni, and S. Schumann, *Eur. Phys. J. C* **71**, 1541 (2011), 0906.2474.
- [29] C. Degrande, C. Duhr, B. Fuks, D. Grellscheid, O. Mattelaer, and T. Reiter, *Comput. Phys. Commun.* **183**, 1201 (2012), 1108.2040.
- [30] L. Darmé et al. (2023), 2304.09883.
- [31] C. Herwig, J. Isaacson, B. Jayatilaka, and P. A. N. Machado, *DarkNeutrino FeynRules Files* (2023), GitHub link: <https://github.com/AchillesGen/Models/tree/main/kinetic-mixing>, URL <https://doi.org/10.5281/zenodo.8277735>.
- [32] A. Atre, T. Han, S. Pascoli, and B. Zhang, *JHEP* **05**, 030 (2009), 0901.3589.
- [33] A. de Gouvêa and A. Kobach, *Phys. Rev. D* **93**, 033005 (2016), 1511.00683.
- [34] R. Aaij et al. (LHCb), *Phys. Rev. Lett.* **124**, 041801 (2020), 1910.06926.
- [35] C. Anastasiou, L. J. Dixon, K. Melnikov, and F. Petriello, *Phys. Rev. D* **69**, 094008 (2004), hep-ph/0312266.
- [36] P. J. Rijken and W. L. van Neerven, *Phys. Rev. D* **51**, 44 (1995), URL <https://link.aps.org/doi/10.1103/PhysRevD.51.44>.
- [37] R. Hamberg, W. van Neerven, and T. Matsuura, *Nuclear Physics B* **359**, 343 (1991), ISSN 0550-3213, URL <https://www.sciencedirect.com/science/article/pii/0550321391900645>.
- [38] W. van Neerven and E. Zijlstra, *Nuclear Physics B* **382**, 11 (1992), ISSN 0550-3213, URL <https://www.sciencedirect.com/science/article/pii/055032139290078P>.
- [39] M. Cacciari, M. Czakon, M. Mangano, A. Mitov, and P. Nason, *Phys. Lett. B* **710**, 612 (2012), 1111.5869.
- [40] M. Czakon, P. Fiedler, and A. Mitov, *Phys. Rev. Lett.* **110**, 252004 (2013), 1303.6254.
- [41] S. Chatrchyan et al. (CMS), *JINST* **9**, P10009 (2014), 1405.6569.
- [42] A. M. Sirunyan et al. (CMS), *JINST* **16**, P05014 (2021), 2012.06888.
- [43] CMS Collaboration, *Recording and reconstructing 10 billion unbiased b hadron decays in CMS* (2019), URL <https://cds.cern.ch/record/2704495>.