

Search for long-lived heavy neutral leptons and Higgs portal scalars decaying in the MicroBooNE detector

P. Abratenko,³⁴ J. Anthony,⁴ L. Arellano,¹⁹ J. Asaadi,³³ A. Ashkenazi,³¹ S. Balasubramanian,¹¹ B. Baller,¹¹ C. Barnes,²¹ G. Barr,²⁴ J. Barrow,^{20,31} V. Basque,¹¹ L. Bathe-Peters,¹³ O. Benevides Rodrigues,³⁰ S. Berkman,¹¹ A. Bhandari,¹⁹ M. Bhattacharya,¹¹ M. Bishai,² A. Blake,¹⁶ B. Bogart,²¹ T. Bolton,¹⁵ J. Y. Book,¹³ L. Camilleri,⁹ D. Caratelli,³ I. Caro Terrazas,⁸ F. Cavanna,¹¹ G. Cerati,¹¹ Y. Chen,²⁷ J. M. Conrad,²⁰ M. Convery,²⁷ L. Cooper-Troendle,³⁷ J. I. Crespo-Anadón,⁵ M. Del Tutto,¹¹ S. R. Dennis,⁴ P. Detje,⁴ A. Devitt,¹⁶ R. Diurba,^{1,22} R. Dorrill,¹⁴ K. Duffy,²⁴ S. Dytman,²⁵ B. Eberly,²⁹ A. Ereditato,¹ J. J. Evans,¹⁹ R. Fine,¹⁷ O. G. Finnerud,¹⁹ W. Foreman,¹⁴ B. T. Fleming,³⁷ N. Foppiani,¹³ D. Franco,³⁷ A. P. Furmanski,²² D. Garcia-Gamez,¹² S. Gardiner,¹¹ G. Ge,⁹ S. Gollapinni,^{32,17} O. Goodwin,¹⁹ E. Gramellini,¹¹ P. Green,¹⁹ H. Greenlee,¹¹ W. Gu,² R. Guenette,^{13,19} P. Guzowski,¹⁹ L. Hagaman,³⁷ O. Hen,²⁰ R. Hicks,¹⁷ C. Hilgenberg,²² G. A. Horton-Smith,¹⁵ R. Itay,²⁷ C. James,¹¹ X. Ji,² L. Jiang,³⁵ J. H. Jo,³⁷ R. A. Johnson,⁷ Y.-J. Jwa,⁹ D. Kalra,⁹ N. Kamp,²⁰ N. Kaneshige,³ G. Karagiorgi,⁹ W. Ketchum,¹¹ M. Kirby,¹¹ T. Kobilarcik,¹¹ I. Kreslo,¹ M. B. Leibovitch,³ I. Lepetic,²⁶ J.-Y. Li,¹⁰ K. Li,³⁷ Y. Li,² K. Lin,¹⁷ B. R. Littlejohn,¹⁴ W. C. Louis,¹⁷ X. Luo,³ K. Manivannan,³⁰ C. Mariani,³⁵ D. Marsden,¹⁹ J. Marshall,³⁶ D. A. Martinez Caicedo,²⁸ K. Mason,³⁴ A. Mastbaum,²⁶ N. McConkey,¹⁹ V. Meddage,¹⁵ K. Miller,⁶ J. Mills,³⁴ K. Mistry,¹⁹ A. Mogan,⁸ T. Mohayai,¹¹ M. Mooney,⁸ A. F. Moor,⁴ C. D. Moore,¹¹ L. Mora Lepin,¹⁹ J. Mousseau,²¹ S. Muller-Ababu,¹ D. Naples,²⁵ A. Navrer-Agasson,¹⁹ N. Nayak,² M. Nebot-Guinot,¹⁰ D. A. Newmark,¹⁷ J. Nowak,¹⁶ M. Nunes,³⁰ N. Oza,¹⁷ O. Palamara,¹¹ N. Pallat,²² V. Paolone,²⁵ A. Papadopoulou,²⁰ V. Papavassiliou,²³ H. B. Parkinson,¹⁰ S. F. Pate,²³ N. Patel,¹⁶ Z. Pavlovic,¹¹ E. Piasetzky,³¹ I. D. Ponce-Pinto,³⁷ S. Prince,¹³ X. Qian,² J. L. Raaf,¹¹ V. Radeka,² A. Rafique,¹⁵ M. Reggiani-Guzzo,¹⁹ L. Ren,²³ L. C. J. Rice,²⁵ L. Rochester,²⁷ J. Rodriguez Rondon,²⁸ M. Rosenberg,²⁵ M. Ross-Lonergan,^{9,17} C. Rudolf von Rohr,¹ G. Scanavini,³⁷ D. W. Schmitz,⁶ A. Schukraft,¹¹ W. Seligman,⁹ M. H. Shaevitz,⁹ R. Sharankova,¹¹ J. Shi,⁴ A. Smith,⁴ E. L. Snider,¹¹ M. Soderberg,³⁰ S. Söldner-Rembold,¹⁹ J. Spitz,²¹ M. Stancari,¹¹ J. St. John,¹¹ T. Strauss,¹¹ S. Sword-Fehlberg,²³ A. M. Szelc,¹⁰ W. Tang,³² N. Taniuchi,⁴ K. Terao,²⁷ C. Thorpe,¹⁶ D. Torbunov,² D. Totani,³ M. Toups,¹¹ Y.-T. Tsai,²⁷ J. Tyler,¹⁵ M. A. Uchida,⁴ T. Usher,²⁷ B. Viren,² M. Weber,¹ H. Wei,¹⁸ A. J. White,³⁷ Z. Williams,³³ S. Wolbers,¹¹ T. Wongjirad,³⁴ M. Wospakrik,¹¹ K. Wresilo,⁴ N. Wright,²⁰ W. Wu,¹¹ E. Yandel,³ T. Yang,¹¹ L. E. Yates,¹¹ H. W. Yu,² G. P. Zeller,¹¹ J. Zennamo,¹¹ and C. Zhang²

(MicroBooNE Collaboration)*

¹Universität Bern, Bern CH-3012, Switzerland

²Brookhaven National Laboratory (BNL), Upton, New York 11973, USA

³University of California, Santa Barbara, California 93106, USA

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

⁵Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid E-28040, Spain

⁶University of Chicago, Chicago, Illinois 60637, USA

⁷University of Cincinnati, Cincinnati, Ohio 45221, USA

⁸Colorado State University, Fort Collins, Colorado 80523, USA

⁹Columbia University, New York, New York 10027, USA

¹⁰University of Edinburgh, Edinburgh EH9 3FD, United Kingdom

¹¹Fermi National Accelerator Laboratory (FNAL), Batavia, Illinois 60510, USA

¹²Universidad de Granada, Granada E-18071, Spain

¹³Harvard University, Cambridge, Massachusetts 02138, USA

¹⁴Illinois Institute of Technology (IIT), Chicago, Illinois 60616, USA

¹⁵Kansas State University (KSU), Manhattan, Kansas 66506, USA

¹⁶Lancaster University, Lancaster LA1 4YW, United Kingdom

¹⁷Los Alamos National Laboratory (LANL), Los Alamos, New Mexico 87545, USA

¹⁸Louisiana State University, Baton Rouge, Louisiana 70803, USA

¹⁹The University of Manchester, Manchester M13 9PL, United Kingdom

²⁰Massachusetts Institute of Technology (MIT), Cambridge, Massachusetts 02139, USA

²¹University of Michigan, Ann Arbor, Michigan 48109, USA

²²University of Minnesota, Minneapolis, Minnesota 55455, USA

²³New Mexico State University (NMSU), Las Cruces, New Mexico 88003, USA

²⁴University of Oxford, Oxford OX1 3RH, United Kingdom

²⁵*University of Pittsburgh, Pittsburgh, Pennsylvania 15260, USA*²⁶*Rutgers University, Piscataway, New Jersey 08854, USA*²⁷*SLAC National Accelerator Laboratory, Menlo Park, California 94025, USA*²⁸*South Dakota School of Mines and Technology (SDSMT), Rapid City, South Dakota 57701, USA*²⁹*University of Southern Maine, Portland, Maine 04104, USA*³⁰*Syracuse University, Syracuse, New York 13244, USA*³¹*Tel Aviv University, Tel Aviv 69978, Israel*³²*University of Tennessee, Knoxville, Tennessee 37996, USA*³³*University of Texas, Arlington, Texas 76019, USA*³⁴*Tufts University, Medford, Massachusetts 02155, USA*³⁵*Center for Neutrino Physics, Virginia Tech, Blacksburg, Virginia 24061, USA*³⁶*University of Warwick, Coventry CV4 7AL, United Kingdom*³⁷*Wright Laboratory, Department of Physics, Yale University, New Haven, Connecticut 06520, USA*

(Received 11 July 2022; accepted 31 October 2022; published 28 November 2022)

We present a search for long-lived Higgs portal scalars (HPS) and heavy neutral leptons (HNL) decaying in the MicroBooNE liquid-argon time projection chamber. The measurement is performed using data collected synchronously with the neutrino beam from Fermilab's Main Injector with a total exposure corresponding to 7.01×10^{20} protons on target. We set upper limits at the 90% confidence level on the mixing parameter $|U_{\mu 4}|^2$ ranging from $|U_{\mu 4}|^2 < 12.9 \times 10^{-8}$ for Majorana HNLs with a mass of $m_{\text{HNL}} = 246$ MeV to $|U_{\mu 4}|^2 < 0.92 \times 10^{-8}$ for $m_{\text{HNL}} = 385$ MeV, assuming $|U_{e 4}|^2 = |U_{\tau 4}|^2 = 0$ and HNL decays into $\mu^\pm \pi^\mp$ pairs. These limits on $|U_{\mu 4}|^2$ represent an order of magnitude improvement in sensitivity compared to the previous MicroBooNE result. We also constrain the scalar-Higgs mixing angle θ by searching for HPS decays into $\mu^+ \mu^-$ final states, excluding a contour in the parameter space with lower bounds of $\theta^2 < 31.3 \times 10^{-9}$ for $m_{\text{HPS}} = 212$ GeV and $\theta^2 < 1.09 \times 10^{-9}$ for $m_{\text{HPS}} = 275$ GeV. These are the first constraints on the scalar-Higgs mixing angle θ from a dedicated experimental search in this mass range.

DOI: [10.1103/PhysRevD.106.092006](https://doi.org/10.1103/PhysRevD.106.092006)

I. INTRODUCTION

The MicroBooNE detector [1] began collecting data in 2015, making it the first fully operational detector of the three liquid-argon time projection chambers comprising the Short-Baseline Neutrino (SBN) program [2] at Fermilab. The MicroBooNE detector was exposed to both the booster neutrino beam (BNB) [3] and the neutrino beam from the main injector (NuMI) [4].

We can use the NuMI beam to study beyond-the-Standard Model phenomena such as the production and decay of heavy neutral leptons (HNL) or Higgs portal scalars (HPS), jointly referred to as long-lived particles (LLPs). In addition, we have also used it to measure electron neutrino cross sections on argon [5,6].

In this paper, we present searches for both types of LLPs originating from kaons decaying at rest in the NuMI hadron absorber, which is located downstream of the MicroBooNE

detector at the end of the NuMI decay pipe. The LLP would travel ≈ 104 m to the MicroBooNE detector where it can be detected through its decay. LLPs produced in the absorber would approach the detector in almost the opposite direction to the vast majority of neutrinos, which originate from near the NuMI beam target (Fig. 1).

Event displays of simulated HNL and HPS decays in the MicroBooNE detector are shown in Fig. 2. The signal topology is characterized by exactly two tracks emerging from a common vertex. Since we search for LLPs produced by two-body decays of kaons at rest, these LLPs have a fixed energy for a given mass. These two properties, the direction and energy of the signal LLPs, help to discriminate them from neutrino and cosmic-ray induced background processes. In addition, the kinematics and topologies of HPS decays to $\mu\mu$ pairs and HNL decays to $\mu\pi$ pairs are similar, which allows us to develop a single signal analysis strategy [7].

The MicroBooNE Collaboration has published upper limits on the production of HNLs decaying to $\mu\pi$ pairs for an exposure of 2.0×10^{20} protons on target (POT) from the BNB, using a dedicated trigger configured to detect HNL decays that occur after the neutrino spill reaches the detector. That search yielded upper limits at the 90% confidence level

*microboone_info@fnal.gov

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




FIG. 1. Momentum vectors of simulated neutrino events from the NuMI target (red) and of HPS decays (black), where the HPS originated from kaons decaying at rest (KDAR) in the hadron absorber. The vectors are shown in the $y-z$ plane of the MicroBooNE coordinate system within the detector's active volume, where y points upward and z points in the nominal BNB beam direction. The vectors start at the vertex location, and their lengths are proportional to the magnitude of the momentum of the neutrino or the HPS. Different momentum scales are used for the display.

(CL) on the element $|U_{\mu 4}|^2$ of the extended Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix $|U_{\mu 4}|^2$ for Dirac and Majorana HNLs in the HNL mass range $260 \leq m_{\text{HNL}} \leq 385$ MeV and assuming $|U_{e 4}|^2 = |U_{\tau 4}|^2 = 0$ [8]. Several other collaborations have also set limits on $|U_{\mu 4}|^2$ [9–16]. Also using a liquid-argon detector exposed to the NuMI beam, the ArgoNeuT Collaboration has published a

search for HNL decays into $\nu\mu^+\mu^-$ final states [17]. They derive limits on the mixing matrix element $|U_{\tau 4}|^2$.

The MicroBooNE Collaboration has also published a search for HPS decaying to e^+e^- pairs assuming the HPS originate from kaons decaying at rest after having been produced at the NuMI absorber [18]. A dataset corresponding to 1.93×10^{20} POT is used to set limits on θ^2 in the range 10^{-6} – 10^{-7} at the 95% CL for the mass range directly below the range considered in this paper ($0 < m_{\text{HPS}} < 211$ MeV). Other direct experimental searches for HPS have been published in Refs. [19–23], and reinterpretations of experimental data as HPS limits have been derived in Refs. [24–27].

II. HEAVY NEUTRAL LEPTONS

The HNL is introduced through an extension of the PMNS matrix by adding a single heavy mass eigenstate that mixes very weakly with the three active neutrino states. This minimal extension adds four parameters to the model comprising the HNL mass m_{HNL} and the elements of the extended PMNS matrix, $|U_{\alpha 4}|^2$ ($\alpha = e, \mu, \tau$). The flavor eigenstates are

$$\nu_\alpha = \sum_i U_{\alpha i} \nu_i + U_{\alpha 4} N, \quad (1)$$

with a heavy neutral lepton state N . The HNL production and decay rates are suppressed by the relevant $|U_{\alpha 4}|^2$ element through mixing-mediated interactions with SM gauge bosons.




FIG. 2. Displays of two simulated signal events, where the LLPs approach from the bottom right of the image and then decay in the detector. The left display shows an HNL of mass $m_{\text{HNL}} = 304$ MeV decaying into a muon (track pointing up and right) and a charged pion (track pointing left) that itself subsequently decays into a muon. The right display shows an HPS of mass $m_{\text{HPS}} = 275$ MeV decaying into a μ^- (the long track going left) and a μ^+ (the shorter track), which quickly decays to a Michel positron. The horizontal direction represents the wires on the collection plane, and the vertical direction represents the electron drift time. Colors represent the amount of charge deposited on wires.




FIG. 3. Production of a Majorana HNL state N via mixing in a K^+ meson decay and its subsequent decay into a $\mu^\mp\pi^\pm$ pair. A Dirac HNL will only decay into $\mu^-\pi^+$ pairs.

HNLs would be produced in the decays of charged kaons and pions originating from the proton interactions on the targets of the BNB or NuMI neutrino beams. If the HNL lifetime is sufficiently long to allow the HNL to reach the MicroBooNE detector, they can decay into Standard Model (SM) particles within the argon volume.

We consider the production channel $K^+ \rightarrow \mu^+N$ with the decay $N \rightarrow \mu^\mp\pi^\pm$ as shown in Fig. 3. The HNL production rate and decay width into $\mu\pi$ are each proportional to $|U_{\mu 4}|^2$, and the total rate is therefore proportional to $|U_{\mu 4}|^4$, assuming $|U_{e 4}|^2 = |U_{\tau 4}|^2 = 0$ [28]. We thus place limits exclusively on the $|U_{\mu 4}|^2$ mixing matrix element. The accessible HNL masses are constrained by the condition $m_K - m_\mu > m_{\text{HNL}} > m_\mu + m_\pi$.

HNL states can include Dirac and Majorana mass terms, where Majorana HNLs would decay in equal numbers into $\mu^+\pi^-$ and $\mu^-\pi^+$ final states. Dirac HNLs from K^+ decays could only decay to the charge combination $\mu^-\pi^+$ to conserve lepton number.

III. HIGGS PORTAL SCALARS

The Higgs portal model [29] is an extension to the SM, where an electrically neutral singlet scalar boson mixes with the Higgs boson with a mixing angle θ . Through this mixing, this HPS boson acquires a coupling to SM fermions via their Yukawa couplings to the Higgs boson, which is proportional to $\sin\theta$. The phenomenology of the Higgs portal model at the SBN program, including the equations describing production and decay of the scalar boson, are discussed in Ref. [30]. At the absorber, the dominant production channel will be through the two-body decay $K^+ \rightarrow \pi^+S$ (where the HPS is denoted by S). The dominant decay mechanism is a penguin diagram with a top quark contributing in the loop (Fig. 4).

The partial width for decays to charged leptons with mass m_ℓ is proportional to m_ℓ^2 [30]. If we assume that there are no new dark sector particles with masses $< m_{\text{HPS}}/2$, the branching fraction into $\mu^+\mu^-$ pair is therefore close to 100% for scalar masses in the range $m_{\mu^+\mu^-} < m_{\text{HPS}} < m_{\pi^0\pi^0}$.




FIG. 4. Production of an HPS boson S via mixing in a kaon meson decay and its subsequent decay. Only K^+ mesons contribute to this analysis, although this decay mode can also occur for neutral or K^- mesons.

The decays into $\pi^0\pi^0$ pairs become accessible for $m_{\text{HPS}} > 269$ MeV and the decays into $\pi^+\pi^-$ pairs become accessible at $m_{\text{HPS}} > 279.1$ MeV.

We do not consider decays to neutral pion pairs in this search, since the detector signature differs significantly from muon and charged pions. The $\pi^+\pi^-$ decay signatures appear very similar to $\mu^+\mu^-$ decays. However, the analysis is not sensitive to HPS decays in this decay channel, as the HPS will decay before it reaches the detector due its short lifetime. We set limits as a function of the mixing angle θ as both HPS production and the decay rate are proportional to θ^2 .

IV. MICROBOONE DETECTOR

The MicroBooNE detector [1] is a liquid-argon time projection chamber (LArTPC) situated at near-ground level at Fermilab. It is located at an angle of 8° relative to the NuMI beamline [4] and at a distance of 680 m from the target. The MicroBooNE LArTPC has an active mass of 85 t of liquid argon, in a volume $2.6 \times 2.3 \times 10.4$ m³ in the x , y , and z coordinates of the MicroBooNE coordinate system.¹

Charged particles produced in neutrino interactions with argon or in decays of LLPs will ionize the argon atoms along their trajectories, producing ionization electrons and scintillation light. An electric field of 273 V/cm causes the electrons to drift toward the anode plane, requiring 2.3 ms to drift across the width of the detector.

The anode plane is oriented perpendicular to the electric field and comprises three planes of sense wires with a spacing of 3 mm between adjacent wires and the same spacing separating the wire planes. Ionization electrons induce a bipolar signal when they pass through the first two

¹The MicroBooNE detector is described by a right-handed coordinate system, where the x axis points along the negative drift direction with the origin located at the anode plane, the y axis pointing vertically upward with the origin at the center of the detector, and the z axis points along the direction of the BNB beam, with the origin at the upstream edge of the detector. The polar angle is defined with respect to the z axis, and the azimuthal angle ϕ is defined with respect to the y axis.




FIG. 5. Illustration of the position of the MicroBooNE detector relative to the NuMI target and absorber.

planes of wires, oriented at $\pm 60^\circ$ with respect to the vertical, before being collected on the third plane with vertically oriented wires producing a unipolar signal.

The waveforms measured by the 8192 wires are digitized in a 4.8 ms readout window. This is longer than the 2.3 ms drift time to allow reconstructing out-of-time cosmic rays. The signal processing on the raw waveforms of the Time Projection Chamber (TPC) includes noise filtering and deconvolution to convert wire signals into hit information [31,32]. Subsequently, individual hits corresponding to a localized energy deposit are extracted for each wire. The combination of timing information and energy deposit contained in each waveform is used to create 2D projective views of the event.

An array of 32 8-in. photomultiplier tubes (PMTs) collects the scintillation light produced by argon ionization. Light flashes are reconstructed from the waveforms of the 32 PMTs. To record an event, the NuMI online trigger requires a scintillation light signal in time with the accelerator beam spill window above a configurable threshold for the number of photoelectrons, which was 9.5 for Run 1 and then lowered to 5.75 during Run 3.

A cosmic ray tagger (CRT) surrounding the cryostat was installed about midway through MicroBooNE operations [33]. It comprises four panels made up of interleaved plastic scintillator strips placed above, below, and on the two sides parallel to the BNB beam direction. The CRT provides both fast timing and positional information of cosmic rays entering the TPC.

V. NUMI BEAMLINE

Protons with an energy of 120 GeV from the Main Injector hit a graphite target, producing particles in the NuMI beamline. The position of the MicroBooNE detector relative to the components of the NuMI beamline is shown in Fig. 5. A system of electromagnetic horns focuses the

charged particles either toward or away from the beam axis, depending on the horn polarity. In forward horn current (FHC) mode, a positive (+200 kA) current is applied to the horns, which focuses positively charged particles in the beam direction. In reverse horn current (RHC) mode, a negative current (−200 kA) is applied to focus negatively charged particles. In this paper, we use NuMI data collected in both modes.

The focused particles then travel down a 675 m long decay pipe filled with helium where they decay to neutrinos or antineutrinos. The proton beam structure determines the intensity and timing structure of the neutrino beam. The neutrino beam has six batches which together form a spill. Each spill is 9.6 μ s long.

Immediately downstream of the decay pipe, at a distance of ≈ 104 m from the MicroBooNE detector, is an absorber (5.5 m wide, 5.6 m tall, and 8.5 m deep) made of an aluminium core surrounded by steel and then concrete, designed to absorb the remaining hadrons. In the MicroBooNE coordinate system, the direction from the absorber corresponds to $\theta = 2.20$ and $\phi = 1.15$ (in radians). Approximately 13% of the beam protons pass through the target without interacting, traveling along the decay pipe before colliding with the absorber at a distance of 725 m downstream from the target [4]. These collisions produce a large number of K^+ mesons which then decay at rest, while most of the K^- mesons are absorbed. The LLPs studied in this paper would be produced in this absorber from K^+ decays. We assume equal rates of K^+ production for the two horn polarities [34].

VI. FLUX GENERATION

The flux of neutrinos in the NuMI beam is simulated in several steps as described in Ref. [5]. The NuMI beamline simulation uses the G4NUMI code [35], which is based on a GEANT4 description of the geometry. The PPFX software

package provides a neutrino flux prediction and uncertainties on the flux [36]. We need to estimate the number of kaons decaying at rest in the absorber, in order to simulate the LLP signal with the G4NUMI beamline simulation.

The NuMI flux from the absorber leads to a negligible rate of neutrino interactions contributing to, e.g., MicroBooNE cross-section results [5,6]. The modeling of the flux of neutrinos from kaons decaying at rest in the absorber is, however, relevant for the LLP signal simulation. We therefore improve on this flux simulation by incorporating previous work from the MiniBooNE Collaboration.

The MiniBooNE detector is located downstream from the MicroBooNE detector in the NuMI beam. The MiniBooNE Collaboration measured the ν_μ flux from kaons decaying at rest at the NuMI absorber [34] using several methods and compared the measurements to several predictions. They adopted the GEANT4 prediction of 0.085 ν_μ produced per POT as their central value, with a 30% uncertainty taken from the range between simulations. The flux predicted by G4NUMI is almost an order of magnitude smaller than the MiniBooNE central value and is not consistent with the 30% uncertainty. The G4NUMI flux would yield a neutrino cross section measurement with the MiniBooNE detector that is inconsistent with expectations. We therefore use the MiniBooNE central flux value for the signal simulation, consistent with the procedure in Ref. [18].

VII. SIGNAL KINEMATICS

We generate the LLP signal using the flux of charged kaons that produce neutrinos, decaying them instead into an HNL or an HPS through the processes $K \rightarrow \mu N$ or $K \rightarrow \pi S$. The two-body decay is isotropic in the kaon's rest frame, and the energy E_{LLP} of the LLP is given by

$$E_{\text{LLP}} = \frac{m_K^2 + m_{\text{LLP}}^2 - m_{\mu,\pi}^2}{2m_K}, \quad (2)$$

where m_{LLP} is the LLP mass, m_K is the kaon mass, and $m_{\mu,\pi}$ is the mass of the daughter particle, which is either a muon or a pion.

The HPS decays into the $\mu^+\mu^-$ final state are simulated to reproduce the branching fraction as a function of HPS mass, with an isotropic decay distribution in the HPS rest frame. We simulate both charge conjugations of HNL decays, i.e., $\mu^-\pi^+$ and $\mu^+\pi^-$ final states, again with isotropic angular distributions. The angular distributions are then reweighted for Dirac and Majorana HNL decays [37].

We assume a uniform time distribution of the proton beam interacting in the target, with a NuMI beam window of 9.6 μs . The two daughter particles are boosted to the laboratory frame using the LLP momentum vector. The time of the LLP decay is calculated from the kaon decay

time and the time of flight of the LLP to the decay point. All times of flight are taken into account in the simulation.

The LLP is only retained if its momentum vector intersects the detector volume. A decay vertex within the detector volume is then chosen along the trajectory of the LLP. The exponential decay of the HPS flux is accounted for when selecting a decay vertex. The HPS lifetime is proportional to θ^{-2} [30]. The decay length of an HPS with $m_{\text{HPS}} > 2m_\mu$ and a mixing angle in the region of interest of $10^{-7} < \theta < 10^{-9}$ is similar to the distance between the absorber and the MicroBooNE detector. Therefore, some HPS will decay before reaching the detector, reducing the flux in the MicroBooNE detector.

For large values of θ^2 , only a small fraction of the HPS reaches the detector before decaying, which restricts the upper reach of exclusion contours as a function of θ^2 . To derive these contours, we first define an effective angle θ_{eff} which neglects the impact of decays before the detector. The event rate in the MicroBooNE detector is then $\propto \theta_{\text{eff}}^4$. In the final extraction of the limits, we correct for the early decays to obtain the limits as a function of the physical mixing angle θ^2 .

The exponential decay of the HNL flux is negligible for the mixing angles considered here as the HNL lifetime is much longer than the time needed to reach the MicroBooNE detector. The number of HNL decaying before reaching the detector is therefore neglected, and the final event rate is proportional to $|U_{\mu 4}|^4$.

VIII. SIMULATION AND RECONSTRUCTION

Simulation and reconstruction are performed within the LARSOFT [38] framework. The neutrino event generator GENIE [39] simulates the neutrino interactions on argon inside the cryostat as well as interactions with the surrounding material. The GENIE configuration used in the simulation is found in Ref. [40], which describes a tuning of phenomenological parameters related to charged-current quasielastic scattering and scattering on a pair of correlated nucleons, based on a fit to external data.

To obtain the response of the detector to ionization charge and scintillation light, the propagation through the detector material of secondary particles produced in the LLP decays or in the neutrino interactions is simulated by GEANT4 [41].

Cosmic rays crossing the detector in the same readout window are taken from data recorded outside of the beam window and then overlaid on the Monte Carlo (MC) simulation. This overlay also addresses the need for simulating detector noise, as the simulated waveforms are combined with waveforms from data recorded during beam-off times.

The simulated samples and the data are reconstructed using the same algorithms. We use the PANDORA pattern recognition framework [42] to combine hits, first clustered independently on each anode plane and then combined

across anode planes to build particles reconstructed in three dimensions. The particles are arranged in a parent-daughter hierarchy based on the topology of the event and classified as tracklike (muons, charged pions, and protons) or shower-like (electrons and photons).

Optical hits are constructed from PMT waveforms. Time-coincident optical hits from different PMTs are combined to form “flashes” that are attributed to a single interaction in the detector. The time of the flash, the associated location, and the amount of light are determined for each flash.

IX. DATASETS

In total, we recorded data corresponding to 2.23×10^{21} POT with the MicroBooNE detector exposed to the NuMI beam. For this paper, we analyze a subset of the data corresponding to 7.01×10^{20} POT, which were taken in two different operating modes, forward horn current and reverse horn current. The FHC dataset was recorded during Run 1 in 2015–2016, and the RHC dataset was recorded during Run 3 in 2017–2018. The CRT was fully installed for the second period, where it is used in the analysis. We thus analyze the two datasets separately to account for differences in neutrino flux, detector configuration, and CRT coverage.

The “beam-on” samples are defined by triggers that are coincident with the neutrino beam; they are used to search for the LLP signal. In addition, we use three samples for each data period (FHC and RHC) that are designed to describe the background (“beam off,” “in-cryostat neutrino,” and “out-of-cryostat neutrino”).

The majority of beam-on events do not contain a neutrino interaction in MicroBooNE and are triggered by a cosmic ray. This source of background is modeled using a sample of events collected under identical trigger conditions but at times where no neutrino beam is present. These samples are referred to as beam off. The beam-off sample is scaled so that its normalization corresponds to the number of beam spills of the beam-on sample. An additional scaling factor of 0.98 is applied to the beam-off sample. This factor takes into account that about 2% of all NuMI spills contains a neutrino interaction in the detector, and the remaining spills contain only cosmic rays [43].

The neutrino-induced background from the NuMI beam is modeled using two samples for each run period. The “in-cryostat ν ” sample contains interactions of neutrinos with the argon inside the cryostat, and the “out-of-cryostat ν ” sample describes interactions with the detector structure and surrounding material. Both samples are generated with the MC simulation and normalized to the number of POT. The number of POT and the number of events passing the online-trigger before scaling factors are applied are summarized in Table I.

TABLE I. Number of events before applying scaling factors for the data and background samples, separated by run periods. The corresponding number of POT are given for the beam-on data and for the simulated background samples (in-cryostat ν and out-of-cryostat ν).

	Run 1 (FHC)		Run 3 (RHC)	
	POT $\times 10^{20}$	Events $\times 10^5$	POT $\times 10^{20}$	Events $\times 10^5$
Data sample:				
Beam on	2.00	6.11	5.01	11.04
Background samples:				
Beam off		9.12		32.37
In-cryostat ν	23.3	9.11	19.8	7.46
Out-of-cryostat ν	16.7	5.69	10.3	3.86

After matching the sample size to the number of POT of the beam-on data, the out-of-cryostat ν samples are scaled by an additional factor to ensure that the sum of the beam-off, in-cryostat, and out-of-cryostat ν samples matches the data normalization within the NuMI timing window of 5.64–15.44 μs , as shown in Fig. 6. There is a small residual difference of $\approx 2\%$ in the normalization of data relative to the sum of the background samples at this stage, as the scaling factors are derived with only a subset of the data.

Simulated datasets using the procedure described in Sec. VIII are used to evaluate the reconstruction and selection efficiency for the signal LLP decays. In total, we




FIG. 6. Distributions of flash times for all Run 3 data events. The beam-on data are compared to the sum of the beam-off, in-cryostat ν , and out-of-cryostat ν samples. The single bin at the start of the neutrino spill at 5.64 μs , where the expectation significantly exceeds the data, is caused by an effect in the timing structure of the beam that is not included in the simulation. The overall normalization of the data and prediction agree within $\approx 2\%$.




FIG. 7. Expected time of arrival of HNLs as a function of m_{HNL} . The horizontal lines show the timing range consistent with the expected time of arrival of neutrinos from the NuMI target. The trigger window extends by about $1 \mu\text{s}$ on both sides compared to the timing range for neutrino interactions.

generate 12 samples in the mass range $246 \leq m_{\text{HNL}} \leq 385 \text{ MeV}$ and eight samples in the mass range $212 \leq m_{\text{HPS}} \leq 279 \text{ MeV}$. The spacing of the mass points takes into account the resolution. Additional points were generated at the edges of phase space, where limits could change more rapidly.

The expected time of arrival of HNLs as a function of m_{HNL} is shown in Fig. 7. The MicroBooNE NuMI trigger requires a flash in a timing window of $4.69\text{--}16.41 \mu\text{s}$ that covers the arrival times of the NuMI neutrino beam at the detector. More than 95% of HNLs arrive within the trigger window at low m_{HNL} . Since the time of flight increases with m_{HNL} , the trigger efficiency decreases to $\approx 55\%$ for $m_{\text{HNL}} = 385 \text{ MeV}$ and then quickly goes to zero. As we study a lower mass range for the HPS, this effect is not relevant for the HPS search.

X. CANDIDATE VERTEX IDENTIFICATION

The Pandora reconstruction algorithm groups objects into “slices” after removing cosmic-ray related hits. The slice identification (SliceID) process uses a combination of charge and light information to identify whether a slice corresponds to a neutrino interaction. Slices containing through-going, out-of-time, or stopping muons are removed. The reconstructed charge in the slice is required to be consistent with the location and intensity of the flash that has triggered the event. If this selection leaves ≥ 2 remaining slices in the event, a Support Vector Machine (SVM) is applied to calculate a “topological score” to select the most “neutrino-like” of the slices. This slice is then also used for the LLP search.

The LLP topology is characterized by a single vertex with two decay particles. The PANDORA reconstruction frequently misplaces the location of the LLP decay vertex as the signal topology differs from the standard neutrino interactions (see Figs. 1 and 2). The vertex locations are therefore recalculated for this analysis. All pairs of fitted tracks whose start or end points lie within a 3D distance of 5 cm of each other are combined to form signal candidates with a common vertex. If the end point of a track is placed at one of the new candidate vertices, the track direction is reversed. The vertex coordinates are calculated as the mean of the start locations of the two tracks. At this stage, there can be multiple vertices in an event. Approximately (60–70)% of signal events contain at least one vertex, and (40–50)% contain exactly one vertex for the generated signal samples (Table II).

XI. KINEMATIC AND TOPOLOGICAL VARIABLES

We define several kinematic and topological variables to discriminate between signal and background. These variables are either associated to the slice or to the signal

TABLE II. Number of events with at least one candidate vertex and the total number of vertices for events that pass the candidate selection, normalized to the number of POT of the beam-on data sample. The number of expected HNL candidates is scaled to $|U_{\mu 4}|^2 = 10^{-8}$, and the number of HPS candidates is scaled to $\theta_{\text{eff}}^2 = 10^{-9}$. The percentages are calculated with respect to the number of events that pass the PANDORA SliceID procedure.

Sample	Run 1 (FHC)		Run 3 (RHC)	
	Events	Vertices	Events	Vertices
Beam off	24033 (38.8%)	37313	50068 (38.6%)	77776
In-cryostat ν	17736 (58.3%)	43226	45582 (59.8%)	119443
Out-of-cryostat ν	3620 (30.1%)	4952	4557 (27.8%)	6418
Sum of predictions	45390 (43.5%)	85490	100207 (45.1%)	203636
Beam on (data)	45638 (44.1%)	86834	98061 (44.1%)	194260
Data/prediction	1.01	1.00	0.98	0.95
Signal				
$m_{\text{HNL}} = 304 \text{ MeV}$	10.7 (64.7%)	16.0	26.5 (64.6%)	40.2
$m_{\text{HPS}} = 245 \text{ MeV}$	8.8 (64.4%)	13.6	22.1 (66.7%)	34.8

candidate. The slice-related variables use information calculated for the entire slice:

- (i) *Multiplicity*: The total multiplicity of objects in the slice, N_{tot} , and the multiplicities of the objects classified as either tracks or showers (N_{tr} or N_{sh}).
- (ii) *Containment*: Events are required to be contained inside the TPC's active volume by restricting the maximum and minimum coordinates of the start and end points of the objects within the slice to be

$$\begin{aligned} 9 < x < 253 \text{ cm}, \\ -112 < y < 112 \text{ cm}, \\ 14 < z < 1020 \text{ cm}. \end{aligned} \quad (3)$$

We also define the maximum and minimum extent of the slice for each coordinate, denoted as $\max(i)$ and $\min(i)$, with $i = x, y, z$. They are defined as the largest or smallest value of the start and end points of tracks in the slice in each dimension.

- (iii) *Slice energy*: The energy of all the objects of the slice, E_{sl} , is reconstructed from the charge readout by the collection planes. We expect that

$$E_{\text{sl}} \approx E_{\text{LLP}} - (m_1 + m_2), \quad (4)$$

where E_{LLP} is the energy of the LLP and m_1 and m_2 are the masses of the two decay particles. This assumes the decay particles are stable and their kinetic energy is fully measured in the TPC. The ionization energy contributed by secondary decays will increase E_{sl} .

- (iv) *Topological score*: A Support Vector Machine (SVM) is applied to calculate a topological score that selects the slice that is most consistent with containing a neutrino interaction.

The remaining variables are directly related to the signal candidate:

- (i) *Proton likelihood*: The log-likelihood particle identification score, S_{PID} , is designed to utilize the calorimetric information to discriminate between protons and minimally ionizing particles (muons) [44]. A score of $S_{\text{PID}} = -1$ indicates that a track is consistent with the proton hypothesis, and $S_{\text{PID}} = +1$ indicates that a track is consistent with the muon hypothesis.
- (ii) *Track length*: We calculate the lengths ℓ_i of the two tracks defining the candidate in the TPC.
- (iii) *Candidate four-momentum*: The momenta of the particles associated to the two tracks produced by the LLP decay are determined using the length of the tracks and the continuous-slowing-down-approximation [45]. For the HNL decays, the momenta and candidate mass are calculated assuming that the longer track is the muon and the shorter track

is the pion. The momenta are summed to calculate the LLP candidate's four-momentum. We define β as the angle between the momentum of the candidate and the vector connecting the center of the absorber to the candidate vertex. The angle α is the opening angle of the two tracks.

XII. EVENT SELECTION

The first stage of the final selection requires that the flash time coincides with the beam window, and only the flash with the largest number of associated photoelectrons is used. If there is a CRT hit within 1 μs of the flash for the Run 3 (RHC) sample, the event is identified as a cosmic ray and rejected. The 3D distance between the Pandora vertex and the closest tracks reconstructed using the CRT is required to be > 20 cm to remove events where charge from cosmic rays not in time with the beam flash is reconstructed as part of a candidate.

We require $2 \leq N_{\text{tot}} \leq 4$, $N_{\text{tr}} \leq 3$, and $N_{\text{sh}} \leq 2$. The event must be contained, and the reconstructed energy $E_{\text{sl}} < 500$ MeV. To remove neutrino interactions producing a proton, candidates where at least one track has a score $S_{\text{PID}} < -0.5$ are removed. Candidates are required to have an opening angle $\cos \alpha > -0.94$. This removes events where a single cosmic ray is split into two tracks with a large opening angle. Finally, the longest track must be shorter than 50 cm.

Table III shows the combined efficiency of the selection requirements and vertex reconstruction relative to the SliceID requirements for background and signal. The background rejection is better for Run 3 compared to Run 1 as the CRT improves the rejection of cosmic rays. The CRT also vetoes particles produced in neutrino interactions that either enter or exit the detector, as expected for out-of-cryostat and in-cryostat neutrino interactions. The effect of the CRT veto

TABLE III. Number of events with at least one candidate vertex that pass in each sample after the full event selection. All numbers are POT normalized. HNL candidates are scaled to $|U_{\mu 4}|^2 = 10^{-8}$, and HPS candidates are scaled to $\theta_{\text{eff}}^2 = 10^{-9}$. The efficiencies are given with respect to the number of events that pass the PANDORA SliceID procedure.

Sample	Run 1 (FHC)		Run 3 (RHC)	
	Events	Efficiency	Events	Efficiency
Beam off	1552	2.5%	1234	0.9%
In-cryostat ν	1188	3.9%	2132	2.7%
Out-of-cryostat ν	208	1.7%	129	0.8%
Sum of predictions	2948	2.8%	3495	1.6%
Beam on (data)	2950	2.8%	3410	1.5%
Data/prediction	1.00		0.99	
Signal				
$m_{\text{HNL}} = 304$ MeV	7.5	45.2%	17.6	43.1%
$m_{\text{HPS}} = 245$ MeV	6.2	45.9%	15.1	45.6%

on the signal is small, as the decay tracks are shorter. The signal efficiency relative to the SliceID requirements are in the range (36–45)%, rising with LLP mass. The total efficiency for an LLP decaying in the detector to be triggered, to be reconstructed and selected by the SliceID, and to pass the final event selection is in the range (13–30)%, again increasing with LLP mass.

XIII. BOOSTED DECISION TREE

We use the XGBOOST gradient boosting library [46] to train boosted decision trees (BDTs) that discriminate between the LLP signal and the background candidates passing the initial selection. A separate BDT is trained for each m_{HNL} and m_{HPS} mass point. The BDTs are trained with 30% of the selected in-cryostat ν sample and 50% of each of the generated signal samples. To improve performance of the BDT training, we require that $> 90\%$ of the hits in the slice are created by the LLP decay products. This procedure excludes misreconstructed signal events.

The background sample contains events where hits from overlaid cosmic events are misreconstructed as signal candidates. Therefore, cosmic-ray background is rejected by the BDT, even without training on a beam-off sample. In total, we use 21 BDT input variables:

- (i) slice energy E_{sl} ;
- (ii) topological score;
- (iii) maximum extent of the slice, $\max(i)$, and minimum extent, $\min(i)$, with $i = x, y, z$;
- (iv) multiplicities N_{tot} , N_{sh} , and N_{tr} ;
- (v) candidate angle β ;
- (vi) candidate opening angle α ;
- (vii) candidate mass m_{LLP} ;
- (viii) angles θ and ϕ of the longest track, as defined in the MicroBooNE coordinate system;
- (ix) length, the PID score S_{PID} and the track score of the two tracks forming the candidate.

To compare data and background simulation, and to demonstrate the sensitivity to an LLP signal, we show eight of the more important variables in the BDT in Fig. 8. In general, we observe good agreement between the background prediction and the data across all variables, both in shape and normalization. The distributions shown are for Run 3 (RHC). The distributions for Run 1 (FHC) are similar with an increased cosmic-ray contribution (see Table III) as the CRT was not yet operational.

In Fig. 8, we overlay the signal distribution for a typical mass point, at $m_{\text{HNL}} = 304$ MeV. As expected, the average reconstructed invariant mass is centered around this value. The slice energy peaks at $E_{\text{sl}} \approx 105$ MeV, about 20 MeV above the expected value from Eq. (4). This shift is related to the ionization energy deposited by Michel electrons. The slice energy is the variable with the best sensitivity to signal in the BDTs. The candidate opening angle is expected to peak at $\cos \alpha = -0.34$ for this mass. The candidate angle β

with respect to the direction from the absorber is peaked at $\cos \beta = 1$ for a large fraction of the HNL signal candidates.

The lengths of the tracks are measures of the momenta of the two particles, which depend on the decay kinematics. The direction of the longer decay track is included as the angles θ and ϕ of the track in the detector coordinate system. The direction (in radians) from the absorber corresponds to $\theta = 2.20$ and $\phi = 1.15$. The angle ϕ helps to reject cosmic rays, which are aligned with $\phi = \pm\pi/2$.

Figure 9 shows the distributions of the BDT scores trained for three different HNL and HPS masses. The BDTs offer strong rejection against background candidates from neutrino interactions. Background candidates from misreconstructed cosmic-ray tracks, as found in the beam off sample, are also rejected with high efficiency. The discrimination between signal and background improves slightly with LLP mass due to the higher energy of the decay particles, leading to improved kinematic reconstruction.

In Fig. 9, the full range of BDT scores lies in the range [0,1]. The final score distribution is shown after a transformation using the inverse of the logistic function has been applied, which maps the score to a range $[-\infty, \infty]$. We only consider candidates with BDT score > 0 after the transformation, corresponding to > 0.5 in the original distribution, since this region contains $> 90\%$ of the signal and hence dominates the sensitivity.

XIV. SYSTEMATIC UNCERTAINTIES

Uncertainty sources are considered for the background samples and signal samples by applying variations that modify the BDT score distributions. For the simulated in-cryostat ν background sample, we consider the impact of the flux simulation, cross section modeling, hadron interactions with argon, and detector variations:

- (i) *Flux simulation.*—Uncertainties on the neutrino flux arise primarily from the rates and kinematics of hadron production in the beamline. The PPFx package is used to estimate these uncertainties. Each of the parameters in the constrained flux prediction is simultaneously sampled within its estimated uncertainties to produce alternative flux predictions. Flux uncertainties also include variations in the beamline conditions, e.g., changes in horn current, horn position, and beam spot location. However, these beamline condition uncertainties are neglected, as they have been shown to be small [5,6] and were therefore not reassessed.
- (ii) *Cross section modeling.*—We assess the uncertainties due to the modeling of neutrino cross sections by varying 44 of the parameters used by the GENIE generator. A full discussion can be found in Ref. [40].
- (iii) *Hadron interactions.*—Hadrons produced in neutrino interactions interact strongly, affecting their propagation through argon. The description of hadron interactions therefore impacts the event




FIG. 8. Sum of the background predictions (beam-off, in-cryostat ν , and out-of-cryostat ν) for several BDT input variables: (a) slice energy E_{sl} , (b) candidate angle β , (c) candidate opening angle α , (d) candidate mass m_{LLP} , (e) length ℓ_t of longer track, (f) length ℓ_t of shorter track, (g) angle ϕ , and (h) angle θ of the longest track in the MicroBooNE coordinate system. The overlaid variable distributions are for an example mass of $m_{\text{HNL}} = 304$ MeV. An arbitrary signal normalization is used for visibility. The distributions are shown for Run 3. The gray bands represent the combined statistical and systematic uncertainty on the prediction.

reconstruction and the description of the neutrino background. These uncertainties are assessed using GEANT4REWEIGHT [47] by considering variations in the GEANT4 cross section model for charged pions and protons.

(iv) *Detector modeling.*—Uncertainties arising from detector modeling are estimated by resimulating the

detector configuration using the same generated input event sample. The variations include effects due to the light simulation, reduction in the light yield, increase of the Rayleigh scattering length, and attenuation of the light in the argon.

Uncertainties related to charge reconstruction are assessed using data-driven modifications of the




FIG. 9. The Run 3 (RHC) BDT performance on the two main background samples and on representative HPS and HNL mass points.

waveforms on the TPC wires, as discussed in Ref. [48]. Additional variations due to the space charge mapping and ion recombination model are simulated and assessed separately.

We expect only a small number of background events in the signal region of the BDT distribution due to the high purity of the event selection. We therefore extrapolate detector-modeling uncertainties from higher statistics regions of the distribution to the signal region, assuming they are constant.

The beam-off sample is taken from data and therefore has no associated systematic uncertainties other than the

statistical fluctuations in the sample. The contribution of the out-of-cryostat sample to the final sample is small. The out-of-cryostat sample normalization (see Sec. IX) is set as an unconstrained parameter in the final fits and found to be negligible.

The impact of the systematic uncertainties on the BDT score distribution is shown in Fig. 10 for the Run 1 (FHC) and Run 3 (RHC) background samples, with the BDT trained for a signal of $m_{\text{HNL}} = 304$ MeV. The fractional uncertainties are given relative to the total background prediction, which is the sum of beam-off, in-cryostat ν , and out-of-cryostat ν samples.




FIG. 10. Fractional uncertainty on the background prediction for the (a) Run 1 (FHC) and (b) Run 3 (RHC) samples as a function of the BDT score, shown separately for the main sources of uncertainty. The BDT has been trained for a signal mass of $m_{\text{HNL}} = 304$ MeV.

The dominant uncertainty in the signal region at high BDT scores is due to the statistical uncertainty of the background samples, which is a consequence of the high purity of the signal selection. Detector modeling uncertainties are $\approx(10-20)\%$, neutrino flux and cross section uncertainties are each $\approx(5-10)\%$ for most bins, while all other uncertainties are negligible. The systematic uncertainties are separately evaluated for all signal training points (signal masses and FHC/RHC), with consistent results. As expected, the sensitivity of the final result is dominated by the RHC dataset.

The dominant contribution to the systematic uncertainty on the signal sample arises from uncertainties on the rate of kaon production at rest in the NuMI absorber. This normalization uncertainty is taken from the evaluation of the MiniBooNE Collaboration to be $\pm 30\%$ [34], as discussed in Sec. VI.

The hadron interaction uncertainty will mainly affect the modeling of pions produced in the HNL decay. The relative uncertainties on the BDT distribution are $\approx 2\%$ for hadron interaction modeling and $< 15\%$ for detector modeling. The longest track is assigned to be the muon in HNL decays; the systematic uncertainty due to this choice is also negligible.

Finally, a 2% uncertainty on the number of POT delivered is estimated from the uncertainty on the beamline toroidal monitor measurement [35], which is taken as fully correlated uncertainty for all samples.

XV. RESULTS

We apply the BDTs trained for each signal mass point to the data and the simulation. In Figs. 11 and 12, we show a comparison of the resulting BDT distributions for the data to the sum of the background predictions, using a selection of representative signal mass points. In events with more than one candidate, we retain only the candidate with the most signal-like BDT score. The background predictions and the data are in good agreement across both of the run periods studied.

The BDT score distributions are used as input to a modified frequentist CL_s calculation [49,50] to set upper limits on the signal strength for each model and mass point. Test statistics are constructed from the log-likelihood ratio (LLR) of the signal-plus-background (S + B) and background-only (B) hypothesis for each LLP mass. The BDT distributions for each run period (FHC and RHC) enter the limit setting as separate channels before their likelihoods are combined.

The systematic uncertainties on the background and signal predictions are taken into account using Gaussian priors. The signal and background predictions are separately fitted to the data distributions under the background-only and signal-plus-background hypotheses. The systematic uncertainties are allowed to vary within their defined priors to maximize the respective likelihood

functions. This reduces the effect of the systematic uncertainties on the sensitivity of the search [51].

The background predictions are fitted separately for the two run periods, with uncorrelated systematic uncertainties. We find that this assumption about the correlations has negligible impact on the result. The flux systematic uncertainties on the signal are taken to be correlated between the two data periods.

The confidence levels are calculated by integrating the LLR distributions, which are derived using pseudoexperiments, under either the S + B (CL_{s+b}) or the B-only hypotheses (CL_b). The excluded signal rate is defined as the scaling of the signal strength for which the confidence level for signal reaches $CL_s = CL_{s+b}/CL_b = 1 - 0.9$.

The observed and median expected 90% CL limits on $|U_{\mu 4}|^2$ are shown for each HNL mass point in Table IV and in Fig. 13, and the corresponding limits on θ_{eff}^2 for the HPS model in Table V. The 1- and 2-standard-deviation intervals cover the range of expected limits produced by 68% and 95% of background prediction outcomes around the median expected value. The observed limits are contained in the 1-standard-deviation interval for all mass points with the exception of $m_{\text{HNL}} = 371.5$ and 385.0 MeV, and for $m_{\text{HPS}} = 215$ MeV, where the observed limit lies within 2 standard deviations. We use a linear interpolation between the mass points when drawing contours.

We derive the HNL limits assuming that HNLs are Majorana particles. For a Dirac HNL, only decays to the charge conjugated final state $\mu^- \pi^+$ are allowed in K^+ decays. The expected number of decays is therefore a factor of 2 smaller for the same $|U_{\mu 4}|^2$ value. The limits for Dirac HNLs are calculated from the Majorana limit by applying a factor of $\sqrt{2}$ to account for the reduced decay rate, since the difference due to angular distributions of the decay is found to be negligible.

Table V shows the values of θ^2 that correspond to the lower and upper bounds of the excluded region for each mass m_{HPS} , where the upper boundary is due to the short lifetime of the HPS. For the highest mass point at $m_{\text{HPS}} = 279$ MeV, the number of HPS reaching the detector before decaying is too small to derive a limit for any of the θ^2 values within the excluded θ_{eff}^2 range.

XVI. COMPARISON WITH EXISTING LIMITS

In Figs. 14 and 15, we compare the observed limits to the existing experimental limits in similar regions of parameter space for both models. The results extend MicroBooNE's sensitivity to $|U_{\mu 4}|^2$ by approximately an order of magnitude compared to the previous MicroBooNE HNL result [8].

Like the MicroBooNE detector, the T2K [9] and NuTeV detectors [10] were located in a neutrino beamline. The PS191 experiment [11,12] at CERN was specifically designed to search for massive decaying neutrinos. The




FIG. 11. Sum of the background predictions (beam-off, in-cryostat ν , and out-of-cryostat ν) for the BDT score distribution compared to data and to the expected signal distribution for $m_{\text{HNL}} = 250, 304,$ and 358 MeV. The HNL signal distributions are normalized to the number of events excluded at the 90% CL (quoted in brackets on the plot), scaled up by a factor of 5 for better visibility and added to the background distribution. The top row shows Run 1 (FHC), and the bottom row shows Run 3 (RHC) data. The bands represent the combined statistical and systematic uncertainty on the prediction.

NA62 [13] and E949 [14] collaborations performed a peak search for HNLs in kaon decays. The muon spectrum measured in stopped $K^+ \rightarrow \mu^+ \nu$ decays ($K_{2\mu}$) has also been used to set limits on HNLs [15,16].

In the mass range $300 < m_{\text{HNL}} < 385$ MeV, this search has similar sensitivity as the NA62 experiment [19]. The E949 [14], PS191 [12], and T2K [9] limits are stronger across the range $300 < m_{\text{HNL}} < 385$ MeV. The T2K Collaboration provides no limit point for masses above 380 MeV. Here, the MicroBooNE limit is of equal or greater sensitivity compared to the NA62 result.

For the HPS model, we constrain a region of parameter space for $212 < m_{\text{HPS}} < 275$ MeV not previously excluded by any dedicated experimental search. The existing limits in this region are reinterpretations of decades

old CHARM [24], LSND [27], and PS191 [26] measurements, performed by authors outside the respective collaborations without access to the original experimental data or MC simulation. Reinterpretations depend on external beamline, flux, and detector simulations. If the signal topology differs from the original selection criteria, the results also depend on estimated detection efficiencies. In the case of the CHARM experiment, the more recent sensitivity estimate in Ref. [24] disagrees by nearly an order of magnitude from the estimate in Ref. [25].

XVII. SUMMARY

We present a search for long-lived particles using NuMI beam data corresponding to 7.01×10^{20} POT recorded with




FIG. 12. Sum of the background predictions (beam-off, in-cryostat ν , and out-of-cryostat ν) for the BDT score distribution compared to data and to the expected signal distribution for $m_{\text{HPS}} = 215, 245,$ and 275 MeV. The HPS signal distributions are normalized to the number of events excluded at the 90% CL (quoted in brackets on the plot), scaled up by a factor of 5 for visibility and added to the background distribution. The top row shows Run 1 (FHC), and the bottom row shows Run 3 (RHC) data. The bands represent the combined statistical and systematic uncertainty on the prediction.




FIG. 13. Limits at the 90% confidence level as function of mass for (a) $|U_{\mu 4}|^2$, assuming a Majorana HNL decaying into $\mu\pi$ pairs, and (b) θ_{eff}^2 of an HPS decaying into $\mu\mu$ pairs. The observed limits are compared to the median expected limit with the 1- and 2-standard-deviation (σ) bands.

TABLE IV. The 90% CL observed and median expected limits on $|U_{\mu 4}|^2$ as a function of m_{HNL} for a Majorana HNL.

m_{HNL} (MeV)	Limit on $ U_{\mu 4} ^2$ ($\times 10^{-8}$)		
	Observed	Median	1 standard deviation range
246.0	12.9	13.7	11.3–17.0
250.0	8.57	7.89	6.46–9.83
263.5	3.86	3.86	3.15–4.71
277.0	3.05	2.55	2.10–3.11
290.5	1.91	1.95	1.59–2.38
304.0	1.46	1.52	1.24–1.85
317.5	1.18	1.19	0.97–1.45
331.0	0.85	0.94	0.77–1.15
344.5	0.67	0.74	0.61–0.92
358.0	0.54	0.65	0.53–0.80
371.5	0.81	0.63	0.51–0.78
385.0	0.92	0.67	0.55–0.83

the MicroBooNE detector. The results are interpreted within two models, where the LLP is either a heavy neutral lepton or a Higgs portal scalar. The LLPs are assumed to be produced by K^+ mesons decaying at rest in the NuMI absorber. The signature in the MicroBooNE liquid-argon detector are HNL decays into $\mu^\pm \pi^\mp$ pairs or HPS decays into $\mu^+ \mu^-$ pairs.

The main sources of background are neutrino and cosmic-ray interactions, where the majority of the neutrino events are from charged-current muon neutrino interactions. To reject background, we select data recorded in time with the NuMI beam and consistent with the LLP signatures in the liquid argon. The LLPs originating in the NuMI absorber enter the detector from a different direction than the majority of beam neutrinos. The decay products also have a fixed energy for a given LLP mass. These kinematic properties are used to discriminate signal from


TABLE V. The 90% CL observed and expected limits on θ_{eff}^2 obtained by this analysis, and the θ^2 contour derived from the θ_{eff}^2 limits.

m_{HPS} (MeV)	Limit on θ_{eff}^2 ($\times 10^{-9}$)			θ^2 range ($\times 10^{-9}$)	
	Observed	Median	1 standard deviation	Low	High
212	30.83	28.25	23.07–34.68	31.3	25000
215	12.1	9.41	7.73–11.5	12.5	1490
230	2.83	2.50	2.04–3.05	3.14	90.0
245	1.36	1.53	1.26–1.86	1.63	32.0
260	1.18	1.01	0.84–1.24	1.55	13.2
269	0.85	0.88	0.72–1.08	1.14	10.2
275	0.82	0.77	0.63–0.95	1.09	5.05
279	0.99	0.80	0.66–0.99		

background. The combined reconstruction and selection efficiencies for LLP decays are between 13% and 30%, increasing with LLP mass. To further improve discrimination between signal and background, we train and apply a BDT with 21 input variables for each mass point.

No significant excess is observed in the BDT score distributions. In the absence of signal, we employ the modified frequentist CL_s method to derive limits on the model mixing parameters $|U_{\mu 4}|^2$ and θ^2 . All limits are presented at the 90% CL.

We set upper limits on the mixing parameter $|U_{\mu 4}|^2$ ranging from $|U_{\mu 4}|^2 = 12.9 \times 10^{-8}$ for Majorana HNLs with a mass of $m_{\text{HNL}} = 246$ MeV to $|U_{\mu 4}|^2 = 0.92 \times 10^{-8}$ for $m_{\text{HNL}} = 385$ MeV, assuming $|U_{e 4}|^2 = |U_{\tau 4}|^2 = 0$ and HNL decays into $\mu^\pm \pi^\mp$ pairs. These limits on $|U_{\mu 4}|^2$ are of similar sensitivity to those published by the NA62 Collaboration [19], and they represent an order of magnitude improvement in sensitivity compared to the previous MicroBooNE result [8].

FIG. 14. Limits on $|U_{\mu 4}|^2$ at the 90% CL as function of mass for Majorana and Dirac HNL compared to the results of the SIN [52], PIENU [53], KEK [16], NA62 [19], E949 [14], PS191 [12], T2K [9], and NuTeV [10] collaborations.




FIG. 15. Limits at the 90% CL on the scalar-Higgs mixing angle θ as a function of m_{HPS} compared to reinterpretations of CHARM [24], LSND [27], and PS191 [26] measurements. In other mass ranges, limits are from a MicroBooNE search for the e^+e^- final state [18] (at the 95% CL) and from searches by the NA62 [19,20] and E949 collaborations [21] for charged kaon decays to pions and an HPS. The LHCb Collaboration performed two searches for an HPS with short lifetime, which would be produced and subsequently decay within the detector [22,23]. The joint coverage of the LHCb result is shown at the 95% CL.

We also constrain the scalar-Higgs mixing angle θ by searching for HPS decays into $\mu^+\mu^-$ final states, excluding a contour in the parameter space with lower bounds of $\theta^2 < 31.3 \times 10^{-9}$ for $m_{\text{HPS}} = 212$ GeV and $\theta^2 < 1.09 \times 10^{-9}$ for $m_{\text{HPS}} = 275$ GeV. These are the first constraints in this region of the θ^2 – m_{HPS} parameter space from a dedicated experimental search. It is also the first search in this mass range using a liquid-argon TPC.

ACKNOWLEDGMENTS

This document was prepared by the MicroBooNE Collaboration using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359. MicroBooNE is supported by the following: the U.S. Department of Energy, Office of Science, Offices of High Energy Physics and Nuclear Physics; the U.S. National Science

Foundation; the Swiss National Science Foundation; the Science and Technology Facilities Council (STFC), part of the United Kingdom Research and Innovation; the Royal Society (United Kingdom); and The European Union’s Horizon 2020 Marie Skłodowska-Curie Actions. Additional support for the laser calibration system and cosmic-ray tagger was provided by the Albert Einstein Center for Fundamental Physics, Bern, Switzerland. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission. We also acknowledge the contributions of technical and scientific staff to the design, construction, and operation of the MicroBooNE detector as well as the contributions of past collaborators to the development of MicroBooNE analyses, without whom this work would not have been possible. We particularly acknowledge the many contributions of Salvatore Davide Porzio to developing this research direction on MicroBooNE.

-
- [1] R. Acciarri *et al.* (MicroBooNE Collaboration), Design and construction of the MicroBooNE detector, *J. Instrum.* **12**, P02017 (2017).
- [2] P. A. Machado, O. Palamara, and D. W. Schmitz, The short-baseline neutrino program at Fermilab, *Annu. Rev. Nucl. Part. Sci.* **69** (2019).
- [3] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), The neutrino flux prediction at MiniBooNE, *Phys. Rev. D* **79**, 072002 (2009).
- [4] P. Adamson *et al.*, The NuMI neutrino beam, *Nucl. Instrum. Methods Phys. Res., Sect. A* **806**, 279 (2016).
- [5] P. Abratenko *et al.* (MicroBooNE Collaboration), Measurement of the flux-averaged inclusive charged-current electron neutrino and antineutrino cross section on argon using the NuMI beam and the MicroBooNE detector, *Phys. Rev. D* **104**, 052002 (2021).
- [6] P. Abratenko *et al.* (MicroBooNE Collaboration), First measurement of inclusive electron-neutrino and antineutrino

- charged current differential cross sections in charged lepton energy on argon in MicroBooNE, *Phys. Rev. D* **105**, L051102 (2022).
- [7] O. Goodwin, Search for Higgs portal scalars and heavy neutral leptons decaying in the MicroBooNE detector, Ph.D. thesis, University of Manchester, 2022.
- [8] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for heavy neutral leptons decaying into muon-pion pairs in the MicroBooNE detector, *Phys. Rev. D* **101**, 052001 (2020).
- [9] K. Abe *et al.* (T2K Collaboration), Search for heavy neutrinos with the T2K near detector ND280, *Phys. Rev. D* **100**, 052006 (2019).
- [10] A. Vaitaitis *et al.* (NuTeV Collaboration), Search for Neutral Heavy Leptons in a High-Energy Neutrino Beam, *Phys. Rev. Lett.* **83**, 4943 (1999).
- [11] G. Bernardi *et al.*, Search for neutrino decay, *Phys. Lett.* **166B**, 479 (1986).
- [12] G. Bernardi *et al.*, Further limits on heavy neutrino couplings, *Phys. Lett. B* **203**, 332 (1988).
- [13] E. Cortina Gil *et al.* (NA62 Collaboration), Search for heavy neutral lepton production in K^+ decays, *Phys. Lett. B* **778**, 137 (2018).
- [14] A. V. Artamonov *et al.* (E949 Collaboration), Search for heavy neutrinos in $K^+ \rightarrow \mu^+ \nu_H$ decays, *Phys. Rev. D* **91**, 052001 (2015); **91**, 059903(E) (2015).
- [15] Y. Asano *et al.*, Search for a heavy neutrino emitted in $K^+ \rightarrow \mu^+ \nu$ decay, *Phys. Lett.* **104B**, 84 (1981).
- [16] R. S. Hayano *et al.*, Heavy Neutrino Search using $K_{\mu 2}$ Decay, *Phys. Rev. Lett.* **49**, 1305 (1982).
- [17] R. Acciarri *et al.* (ArgoNeuT Collaboration), New Constraints on tau-coupled Heavy Neutral Leptons with masses $m_N = 280\text{--}970$ MeV, *Phys. Rev. Lett.* **127**, 121801 (2021).
- [18] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for a Higgs Portal Scalar Decaying to Electron-Positron Pairs in the MicroBooNE Detector, *Phys. Rev. Lett.* **127**, 151803 (2021).
- [19] E. Cortina Gil *et al.* (NA62 Collaboration), Search for K^+ decays to a muon and invisible particles, *Phys. Lett. B* **816**, 136259 (2021).
- [20] E. Cortina Gil *et al.* (NA62 Collaboration), Measurement of the very rare $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decay, *J. High Energy Phys.* **06** (2021) 093.
- [21] A. V. Artamonov *et al.* (BNL-E949 Collaboration), Study of the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ in the momentum region $140 < P_\pi < 199$ MeV/ c , *Phys. Rev. D* **79**, 092004 (2009).
- [22] R. Aaij *et al.* (LHCb Collaboration), Search for Hidden-Sector Bosons in $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ Decays, *Phys. Rev. Lett.* **115**, 161802 (2015).
- [23] R. Aaij *et al.* (LHCb Collaboration), Search for long-lived scalar particles in $B^+ \rightarrow K^+ \chi(\mu^+ \mu^-)$ decays, *Phys. Rev. D* **95**, 071101 (2017).
- [24] M. W. Winkler, Decay and detection of a light scalar boson mixing with the Higgs boson, *Phys. Rev. D* **99**, 015018 (2019).
- [25] J. D. Clarke, R. Foot, and R. R. Volkas, Phenomenology of a very light scalar ($100 \text{ MeV} < m_h < 10 \text{ GeV}$) mixing with the SM Higgs, *J. High Energy Phys.* **02** (2014) 123.
- [26] D. Gorbunov, I. Krasnov, and S. Suvorov, Constraints on light scalars from PS191 results, *Phys. Lett. B* **820**, 136524 (2021).
- [27] S. Foroughi-Abari and A. Ritz, LSND constraints on the Higgs portal, *Phys. Rev. D* **102**, 035015 (2020).
- [28] A. Atre, T. Han, S. Pascoli, and B. Zhang, The search for heavy Majorana neutrinos, *J. High Energy Phys.* **05** (2009) 030.
- [29] B. Patt and F. Wilczek, Higgs-field portal into hidden sectors, [arXiv:hep-ph/0605188](https://arxiv.org/abs/hep-ph/0605188).
- [30] B. Batell, J. Berger, and A. Ismail, Probing the Higgs portal at the Fermilab short-baseline neutrino experiments, *Phys. Rev. D* **100**, 115039 (2019).
- [31] R. Acciarri *et al.* (MicroBooNE Collaboration), Noise characterization and filtering in the MicroBooNE liquid argon TPC, *J. Instrum.* **12**, P08003 (2017).
- [32] C. Adams *et al.* (MicroBooNE Collaboration), Ionization electron signal processing in single phase LArTPCs. part II. data/simulation comparison and performance in MicroBooNE, *J. Instrum.* **13**, P07007 (2018).
- [33] C. Adams *et al.* (MicroBooNE Collaboration), Design and construction of the MicroBooNE Cosmic Ray Tagger system, *J. Instrum.* **14**, P04004 (2019).
- [34] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), First Measurement of Monoenergetic Muon Neutrino Charged Current Interactions, *Phys. Rev. Lett.* **120**, 141802 (2018).
- [35] L. Aliaga *et al.* (MINERvA Collaboration), Neutrino flux predictions for the NuMI beam, *Phys. Rev. D* **94**, 092005 (2016); **95**, 039903(A) (2017).
- [36] L. Aliaga *et al.* (MINERvA Collaboration), Neutrino flux predictions for the NuMI beam, *Phys. Rev. D* **94**, 092005 (2016); **95**, 039903(A) (2017).
- [37] A. B. Balantekin, A. de Gouvêa, and B. Kayser, Addressing the Majorana vs. Dirac question with neutrino decays, *Phys. Lett. B* **789**, 488 (2019).
- [38] E. L. Snider and G. Petrillo, LArSoft: Toolkit for simulation, reconstruction and analysis of liquid argon TPC neutrino detectors, *J. Phys. Conf. Ser.* **898**, 042057 (2017).
- [39] C. Andreopoulos *et al.*, The GENIE neutrino Monte Carlo generator, *Nucl. Instrum. Methods Phys. Res., Sect. A* **614**, 87 (2010).
- [40] P. Abratenko *et al.* (MicroBooNE Collaboration), New CC0 π GENIE model tune for MicroBooNE, *Phys. Rev. D* **105**, 072001 (2022).
- [41] S. Agostinelli *et al.* (GEANT4 Collaboration), GEANT4: A simulation toolkit, *Nucl. Instrum. Methods Phys. Res., Sect. A* **506**, 250 (2003); we use GEANT4_10_1_P03D.
- [42] J. S. Marshall and M. A. Thomson, The pandora software development kit for pattern recognition, *Eur. Phys. J. C* **75**, 439 (2015), we use Pandora v03_00_00.
- [43] K. Mistry, First measurement of the flux-averaged differential charged-current electron-neutrino and antineutrino cross section on argon with the MicroBooNE detector, Ph.D. thesis, University of Manchester, 2021.
- [44] P. Abratenko *et al.* (MicroBooNE Collaboration), Calorimetric classification of track-like signatures in liquid argon

- TPCs using MicroBooNE data, *J. High Energy Phys.* **12** (2021) 153.
- [45] D.E. Groom, N.V. Mokhov, and S.I. Striganov, Muon stopping-power and range tables: 10 MeV–100 TeV, *At. Data Nucl. Data Tables* **78** (1999).
- [46] T. Chen and C. Guestrin, XGBoost: A scalable tree boosting system, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16 (ACM, New York, NY, USA, 2016), pp. 785–794.
- [47] J. Calcutt, C. Thorpe, K. Mahn, and L. Fields, Geant4Re-weight: A framework for evaluating and propagating hadronic interaction uncertainties in Geant4, *J. Instrum.* **16**, P08042 (2021).
- [48] P. Abratenko *et al.* (MicroBooNE Collaboration), Novel approach for evaluating detector-related uncertainties in a LArTPC using MicroBooNE data, *Eur. Phys. J. C* **82**, 454 (2022).
- [49] T. Junk, Confidence level computation for combining searches with small statistics, *Nucl. Instrum. Methods Phys. Res., Sect. A* **434**, 435 (1999).
- [50] A.L. Read, Presentation of search results: The CL_s technique, *J. Phys. G* **28**, 2693 (2002).
- [51] W. Fisher, Systematics and limit calculations, Report No. Fermilab-TM-2386-E, 2006.
- [52] M. Daum, P.-R. Kettle, B. Jost, R.M. Marshall, R.C. Minehart, W.A. Stephens, and K.O.H. Ziock, Search for admixtures of massive neutrinos in the decay $\pi^+ \rightarrow \mu^+ \nu$, *Phys. Rev. D* **36**, 2624 (1987).
- [53] A. Aguilar-Arevalo *et al.* (PIENU Collaboration), Search for heavy neutrinos in $\pi \rightarrow \mu\nu$ decay, *Phys. Lett. B* **798**, 134980 (2019).