

First Search for Neutral Current Coherent Single-Photon Production in MicroBooNE

P. Abratenko,³⁸ D. Andrade Aldana,¹⁴ L. Arellano,²¹ J. Asaadi,³⁷ A. Ashkenazi,³⁶ S. Balasubramanian,¹² B. Baller,¹² A. Barnard,²⁸ G. Barr,²⁸ D. Barrow,²⁸ J. Barrow,²⁵ V. Basque,¹² J. Bateman,^{15,21} O. Benevides Rodrigues,¹⁴ S. Berkman,²⁴ A. Bhat,⁷ M. Bhattacharya,¹² M. Bishai,³ A. Blake,¹⁸ B. Bogart,²³ T. Bolton,¹⁷ M. B. Brunetti,⁴⁰ L. Camilleri,¹⁰ D. Caratelli,⁴ F. Cavanna,¹² G. Cerati,¹² A. Chappell,⁴⁰ Y. Chen,³² J. M. Conrad,²² M. Convery,³² L. Cooper-Troendle,²⁹ J. I. Crespo-Anadón,⁶ R. Cross,⁴⁰ M. Del Tutto,¹² S. R. Dennis,⁵ P. Detje,⁵ R. Diurba,² Z. Djurcic,¹ K. Duffy,²⁸ S. Dytman,²⁹ B. Eberly,³⁴ P. Englezos,³¹ A. Ereditato,^{7,12} J. J. Evans,²¹ C. Fang,⁴ W. Foreman,^{14,19} B. T. Fleming,⁷ D. Franco,⁷ A. P. Furmanski,²⁵ F. Gao,⁴ D. Garcia-Gamez,¹³ S. Gardiner,¹² G. Ge,¹⁰ S. Gollapinni,¹⁹ E. Gramellini,²¹ P. Green,²⁸ H. Greenlee,¹² L. Gu,¹⁸ W. Gu,³ R. Guenette,²¹ P. Guzowski,²¹ L. Hagaman,⁷ M. D. Handley,⁵ O. Hen,²² C. Hilgenberg,²⁵ G. A. Horton-Smith,¹⁷ A. Hussain,¹⁷ B. Irwin,²⁵ M. S. Ismail,²⁹ C. James,¹² X. Ji,²⁶ J. H. Jo,³ R. A. Johnson,⁸ Y.-J. Jwa,¹⁰ D. Kalra,¹⁰ G. Karagiorgi,¹⁰ W. Ketchum,¹² M. Kirby,³ T. Kobilarcik,¹² N. Lane,^{15,21} J.-Y. Li,¹¹ Y. Li,³ K. Lin,³¹ B. R. Littlejohn,¹⁴ L. Liu,¹² W. C. Louis,¹⁹ X. Luo,⁴ T. Mahmud,¹⁸ C. Mariani,³⁹ D. Marsden,²¹ J. Marshall,⁴⁰ N. Martinez,¹⁷ D. A. Martinez Caicedo,³³ S. Martynenko,³ A. Mastbaum,³¹ I. Mawby,¹⁸ N. McConkey,³⁰ L. Mellet,²⁴ J. Mendez,²⁰ J. Micallef,^{22,38} A. Mogan,⁹ T. Mohayai,¹⁶ M. Mooney,⁹ A. F. Moor,⁵ C. D. Moore,¹² L. Mora Lepin,²¹ M. M. Moudgalya,²¹ S. Mullerlababu,² D. Naples,²⁹ A. Navrer-Agasson,^{15,21} N. Nayak,³ M. Nebot-Guinot,¹¹ C. Nguyen,³¹ J. Nowak,¹⁸ N. Oza,¹⁰ O. Palamara,¹² N. Pallat,²⁵ V. Paolone,²⁹ A. Papadopoulou,¹ V. Papavassiliou,²⁷ H. B. Parkinson,¹¹ S. F. Pate,²⁷ N. Patel,¹⁸ Z. Pavlovic,¹² E. Piasetzky,³⁶ K. Pletcher,²⁴ I. Pophale,¹⁸ X. Qian,³ J. L. Raaf,¹² V. Radeka,³ A. Rafique,¹ M. Reggiani-Guzzo,¹¹ J. Rodriguez Rondon,³³ M. Rosenberg,³⁸ M. Ross-Lonergan,¹⁹ I. Safa,¹⁰ D. W. Schmitz,⁷ A. Schukraft,¹² W. Seligman,¹⁰ M. H. Shaevitz,¹⁰ R. Sharankova,¹² J. Shi,⁵ E. L. Snider,¹² M. Soderberg,³⁵ S. Söldner-Rembold,^{15,21} J. Spitz,²³ M. Stancari,¹² J. St. John,¹² T. Strauss,¹² A. M. Szec,¹¹ N. Taniuchi,⁵ K. Terao,³² C. Thorpe,²¹ D. Torbunov,³ D. Totani,⁴ M. Touns,¹² A. Trettin,²¹ Y.-T. Tsai,³² J. Tyler,¹⁷ M. A. Uchida,⁵ T. Usher,³² B. Viren,³ J. Wang,²⁶ M. Weber,² H. Wei,²⁰ A. J. White,⁷ S. Wolbers,¹² T. Wongjirad,³⁸ M. Wospakrik,¹² K. Wresilo,⁵ W. Wu,²⁹ E. Yandel,^{4,19} T. Yang,¹² L. E. Yates,¹² H. W. Yu,³ G. P. Zeller,¹² J. Zennamo,¹² and C. Zhang³

(The MicroBooNE Collaboration)*

¹Argonne National Laboratory (ANL), Lemont, IL, 60439, USA

²Universität Bern, Bern CH-3012, Switzerland

³Brookhaven National Laboratory (BNL), Upton, NY, 11973, USA

⁴University of California, Santa Barbara, CA, 93106, USA

⁵University of Cambridge, Cambridge CB3 0HE, United Kingdom

⁶Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid E-28040, Spain

⁷University of Chicago, Chicago, IL, 60637, USA

⁸University of Cincinnati, Cincinnati, OH, 45221, USA

⁹Colorado State University, Fort Collins, CO, 80523, USA

¹⁰Columbia University, New York, NY, 10027, USA

¹¹University of Edinburgh, Edinburgh EH9 3FD, United Kingdom

¹²Fermi National Accelerator Laboratory (FNAL), Batavia, IL 60510, USA

¹³Universidad de Granada, Granada E-18071, Spain

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

¹⁵Imperial College London, London SW7 2AZ, United Kingdom

¹⁶Indiana University, Bloomington, IN 47405, USA

¹⁷Kansas State University (KSU), Manhattan, KS, 66506, USA

¹⁸Lancaster University, Lancaster LA1 4YW, United Kingdom

¹⁹Los Alamos National Laboratory (LANL), Los Alamos, NM, 87545, USA

²⁰Louisiana State University, Baton Rouge, LA, 70803, USA

²¹The University of Manchester, Manchester M13 9PL, United Kingdom

²²Massachusetts Institute of Technology (MIT), Cambridge, MA, 02139, USA

²³University of Michigan, Ann Arbor, MI, 48109, USA

²⁴Michigan State University, East Lansing, MI 48824, USA

²⁵University of Minnesota, Minneapolis, MN, 55455, USA

²⁶Nankai University, Nankai District, Tianjin 300071, China

²⁷New Mexico State University (NMSU), Las Cruces, NM, 88003, USA

²⁸University of Oxford, Oxford OX1 3RH, United Kingdom

²⁹University of Pittsburgh, Pittsburgh, PA, 15260, USA

³⁰Queen Mary University of London, London E1 4NS, United Kingdom

³¹Rutgers University, Piscataway, NJ, 08854, USA

³²*SLAC National Accelerator Laboratory, Menlo Park, CA, 94025, USA*

³³*South Dakota School of Mines and Technology (SDSMT), Rapid City, SD, 57701, USA*

³⁴*University of Southern Maine, Portland, ME, 04104, USA*

³⁵*Syracuse University, Syracuse, NY, 13244, USA*

³⁶*Tel Aviv University, Tel Aviv, Israel, 69978*

³⁷*University of Texas, Arlington, TX, 76019, USA*

³⁸*Tufts University, Medford, MA, 02155, USA*

³⁹*Center for Neutrino Physics, Virginia Tech, Blacksburg, VA, 24061, USA*

⁴⁰*University of Warwick, Coventry CV4 7AL, United Kingdom*

(Dated: February 11, 2025)

This article presents the first search for neutrino-induced neutral current coherent single-photon production (NC coherent 1γ). The search makes use of data from the MicroBooNE 85-tonne active volume liquid argon time projection chamber detector, situated in the Fermilab Booster Neutrino Beam (BNB), with an average neutrino energy of $\langle E_\nu \rangle \sim 0.8$ GeV. A targeted selection of candidate neutrino interactions with a single photon-like electromagnetic shower in the final state and no visible vertex activity was developed to search for the NC coherent 1γ process, along with two auxiliary selections used to constrain the dominant background from NC π^0 production. With an integrated exposure of 6.87×10^{20} protons on target delivered by the BNB, we set the world's first limit for this rare process, corresponding to an upper limit on the flux-averaged cross section of $\sigma < 1.49 \times 10^{-41} \text{ cm}^2$ at 90% C.L.

I. INTRODUCTION

Neutrino-nucleus cross sections have been the subject of intense study in recent years [1–8] due to their role in interpreting neutrino oscillation measurements and searches for other rare processes in neutrino experiments. Inclusive and exclusive neutral current (NC) neutrino interactions, in particular, are important to understand in their own right and are a contributing background to neutrino oscillation measurements in current and next-generation experiments. NC neutrino interactions are particularly important for ν_e and $\bar{\nu}_e$ measurements in the energy range of less than 1 GeV, especially for detectors that lack robust methods to differentiate between photon- and electron-induced electromagnetic showers. In such detectors, any NC process involving one or more photons in the final state could potentially contribute as a background to ν_e or $\bar{\nu}_e$ charged current (CC) scattering. This necessitates experimental measurements of NC single-photon production to validate or improve existing theoretical predictions for such processes. Misidentification of photons as electrons also complicates the interpretation of ν_e appearance measurements aiming to measure subtle oscillation effects. These include, for example, past short-baseline appearance searches [9, 10], upcoming sterile neutrino oscillation searches with the Short Baseline Neutrino (SBN) program [11], and planned CP violation measurements and mass hierarchy determination with the future Deep Underground Neutrino Experiment (DUNE) [12].

Neutrino-induced NC coherent single-photon production is predicted to be a sub-dominant source of single photons in neutrino-argon scattering below 1 GeV [13]. The dominant single-photon production process in this

energy range is that of NC resonance production of Δ baryons followed by Δ radiative decay [14]. The NC Δ radiative decay signal is expected to be roughly an order of magnitude larger than NC coherent single-photon production on argon. Although this NC coherent single-photon (NC coherent 1γ) process is predicted in the Standard Model (SM), it has never been directly searched for or observed in neutrino scattering.

In this article, we present the first experimental search for neutrino-induced NC coherent 1γ production on argon. This measurement is performed at a mean neutrino energy of $\langle E_\nu \rangle \sim 0.8$ GeV making it especially relevant to the physics programs of experiments operating in a similar energy regime, or with argon as a target material. Additionally, much like MicroBooNE's previous search for NC Δ baryon production and radiative decay [15], this measurement serves as another test for photon [16, 17] and photon-like e^+e^- [18, 19] final state interpretations of the MiniBooNE “low energy excess” [9]. This analysis and the techniques involved also serve as a blueprint for future higher-sensitivity searches for this process that will be enabled, for example, by the upcoming Short Baseline Near Detector (SBND) due to its much higher anticipated event rate [20]. This result is one of three complementary next-generation single-photon searches released simultaneously by MicroBooNE, alongside an updated analysis of the previous NC Δ radiative decay search [21] and an inclusive search for anomalous single-photon events [22].

The interaction final state targeted in this analysis is defined as:

$$\nu(\bar{\nu}) + \text{Ar}_{gs} \rightarrow \nu(\bar{\nu}) + \text{Ar}_{gs} + \gamma, \quad (1)$$

where Ar_{gs} represents the struck and residual (argon) nucleus, which remains in the ground state after scattering. Due to the coherent (low momentum transfer) nature of the interaction, the outgoing photon almost always has

* microboone_info@fnal.gov

a forward direction relative to the beam. Thus, we aim to identify this process in the MicroBooNE detector by searching for a single photon shower in the forward direction. No other observable activity is expected in the final state, and this informs the signal selection strategy. This result makes extensive use of the capability of liquid argon time projection chamber [23, 24] (LArTPC) detectors to identify MeV-scale energy deposits (often called “blips”) in the detector from small, unresolved protons in order to study the coherent nature of selected events in more detail. While there has been much development and demonstration [25, 26] of the potential and capability of low-energy blips to help with neutron identification, supernova and solar neutrino energy reconstruction [27, 28], and ambient background radiation [29], this result represents the first development and application of MeV-scale blip reconstruction in a high-level MicroBooNE physics analysis.

In the remainder of this article we introduce the MicroBooNE experiment in more detail, followed by an overview of the analysis, detailed descriptions of the event selection, systematic uncertainty evaluation, and the sample utilized to constrain the NC π^0 backgrounds in the signal region. We then present the results of our search and conclude with a summary and future outlook.

II. THE MICROBOONE EXPERIMENT

A. The Detector and Neutrino Beam

The MicroBooNE experiment is a LArTPC neutrino experiment located at Fermilab. The MicroBooNE LArTPC detector [30] has a rectangular prism shape with dimensions of 2.3 m in height, 2.6 m in width, and 10.4 m in length, containing ~ 85 metric tons of liquid argon in its active volume. The cathode plane and anode planes are oriented vertically and parallel to the neutrino beam direction, with -70 kV high voltage applied to the cathode plane during normal operation, creating a uniform electric field of 273 V/cm inside the detector. Charged particles traversing the detector ionize the argon atoms, generating ionization electrons along their path. Under the effect of the electric field, these ionization electrons drift toward the anode plane and induce electronic signals. Three sensing wire planes are placed 3 mm apart at the anode; two wire planes record induction signals and are referred to as “plane 0” and “plane 1” in the rest of the paper, and a third wire plane collects the drifting electron charge and is referred to as the collection plane (“plane 2”). The wires within each wire plane are spaced 3 mm apart, and each plane’s wires are oriented at an angle relative to vertical: 0° (“plane 2”), and $\pm 60^\circ$ (“plane 1” and “plane 0”, respectively). Together, these produce three fine-grained two-dimensional (2D) views of the neutrino interaction.

A light collection system consisting of 32 photomultiplier tubes (PMTs) is distributed behind the anode plane

outside of the drift volume. The PMTs collect liquid argon scintillation light produced by the interactions with a time resolution of a few ns. This provides precise timing information for the neutrino interaction and an estimate of the interaction position within the detector to enable cosmic-ray rejection. During data collection, a coincidence requirement between the light collected and the neutrino beam time is applied to reduce cosmic ray backgrounds in the recorded data. In addition, at the reconstruction stage, the magnitude and time of the light signal recorded by each PMT allows time-matching of the light pulses to individual particle candidates in the TPC, resulting in further rejection of cosmic rays. Both the LArTPC and light collection system are immersed in a cylindrical, foam-insulated cryostat vessel holding a total of 170 tonnes of liquid argon and serving as a Faraday cage.

The MicroBooNE detector is exposed to an on-axis neutrino flux from the Booster Neutrino Beam (BNB) [31]. The BNB uses protons accelerated to a momentum of 8.9 GeV/c in the Booster synchrotron and directed onto a beryllium target producing secondary hadrons, such as pions and kaons. These are focused by an electromagnetic horn before decaying in the decay pipe, producing a directed beam of neutrinos. The simulation of the BNB neutrino flux is based on the work by the MiniBooNE Collaboration [31] and adjusted for the MicroBooNE location 72 m upstream of MiniBooNE. The resulting neutrino beam has an average energy of $\langle E_\nu \rangle \sim 0.8$ GeV, consisting predominantly of ν_μ s (93.6%) with a small contamination of $\bar{\nu}_\mu$ (5.8%) and $\nu_e/\bar{\nu}_e$ ($<0.6\%$) [31].

The MicroBooNE detector ran for five years between 2015-2020 and collected approximately 1.3×10^{21} protons-on-target (POT) of BNB data. This analysis makes use of data collected during the first three years of operation (Runs 1-3) from February 2016 to July 2018, corresponding to a total of 6.87×10^{20} POT.

B. Event Simulation and Reconstruction

MicroBooNE utilizes the GENIE [32] neutrino event generator to simulate the neutrino-argon interactions inside and outside the detector. More specifically, GENIE v3.0.6 G18_10a_02_11a is adopted as the default event generator with a customized tune [33] of CC interaction models based on fits to T2K data [34]. As the focus of this analysis is aimed at neutral current interactions, the simulation of background NC π^0 events is of utmost importance. In our GENIE simulation, both coherent and resonant pion production use the Berger-Sehgal model [35], with axial form factors derived from a fit of MiniBooNE data and vector form factors derived from electron scattering fits of BBBA07 [36]. After the neutrino interaction, typically the struck nucleus is left in a highly excited state and will de-excite by emitting nuclear fission fragments, nucleons, or photons. The simulation

of emitted nucleons is via the INTRANUKE [37] intranuclear cascade simulation framework with an empirical, data-driven hA2018 final-state interaction model. Current versions of GENIE only simulate deexcitation photon emission for oxygen, and so these are not modeled in this analysis.

The outgoing final-state particles from GENIE are passed to the LArSoft [38] framework for a full event simulation. LArSoft is a software toolkit that supports event simulation, reconstruction, and analysis for LArTPC experiments. GEANT4 [39] v4_10_3_03c simulates the propagation of particles in the detector, including the energy deposition and production of optical light and ionization electrons along the particle trajectory. The simulation of the propagation of light and ionization charge in the detector is done in LArSoft. Dedicated detector simulation tools [40, 41] are developed to simulate the response of the sensing wires and TPC readout electronics while taking into account different detector effects [42–44].

As a surface detector, MicroBooNE is exposed to a constant cosmic-ray rate of approximately 4.5 kHz [45]. To model the cosmic-ray activity in the detector during the few-millisecond readout window, MicroBooNE takes a data-driven approach. Simulated neutrino event TPC waveforms are overlaid onto data events collected without a PMT optical trigger requirement when the BNB beam was off, to account for cosmic activity during the TPC readout window. The resulting simulated samples are referred to as “overlay” samples. Cosmic-ray backgrounds (“cosmics”) can additionally generate light and charge mimicking a neutrino interaction even when no neutrino interaction happens in the detector during the beam-spill. To evaluate this background, data triggered by the PMTs during beam-off periods is recorded.

The event reconstruction chain in MicroBooNE starts with signal processing. The TPC wire signals are processed through an offline noise filter to remove excessive coherent and incoherent noise from the electronics [40], and then run through 2D wire-time deconvolution to deconvolve the detector response and reconstruct the actual ionization electron distribution [41]. Regions of interest are identified on the deconvolved signals, and Gaussian pulses are extracted by a hit-finding algorithm; the reconstructed Gaussian hits are fed into a reconstruction framework to form high-level reconstruction objects.

This analysis utilizes the Pandora Software Development Kit [46] for pattern recognition and particle trajectory reconstruction. Pandora takes reconstructed Gaussian hits with 2D (time and wire) information as inputs and then uses the temporal and spatial features to perform three-dimensional (3D) reconstruction of the particles. First, Gaussian hits are clustered to form 2D clusters on each plane. A three-dimensional (3D) interaction vertex is then identified from pairs of 2D clusters from different planes while taking into account the distribution of hits of the clusters. Each 2D cluster is assigned a score from 0 to 1 based on the linearity, hit dispersion, extent,

and orientation of the clusters with respect to the reconstructed vertex and assumed beam direction. Next, clusters are classified as “shower-like” or “track-like”. The 3D shower (track) objects are formed by matching 2D shower (track) clusters from the three planes.

Analysis-specific high-level reconstruction is performed on the outputs of the Pandora pattern recognition and particle trajectory reconstruction. Due to the single-shower topology of the targeted signal, we limit the discussion here to the high-level reconstruction only of the reconstructed shower. The dE/dx of the shower start is evaluated by identifying the main trunk of the shower start through a Kalman filter based procedure [47] and taking the median of the dE/dx calculated over the first 4 cm of the start of shower. The energy of the shower is calculated by summing up the energy of all hits associated with the shower on a plane-by-plane basis and taking the maximum of all three planes. The resulting calorimetric energy of the shower is found to be systematically about 20% lower than the true shower energy in simulated events [48]. This is due to lossy effects such as thresholding, where small energy depositions during the shower cascade are too low to be reconstructed, and misclustering, where reconstructed hits from the shower activity are missed during clustering. We account for this by correcting the reconstructed shower energy following the procedure in Ref. [49] where an approximate 20% increase is applied to ensure that the π^0 mass peak reconstructed from our Monte-Carlo (MC) simulations agrees with the measured value.

III. ANALYSIS OVERVIEW

This analysis searches for the NC coherent 1γ production on argon. The theoretical model is proposed and described in detail in [13], with a predicted total cross section on argon of $\frac{\sigma}{A} \sim \mathcal{O}(10^{-43} \text{ cm}^2)$ where A is the atomic mass number of the target nucleus. This model is not implemented in the GENIE v3.0.6 generator used by the default MicroBooNE software; instead it has been implemented and validated in an updated GENIE version v.3.2 [50] on the development branch tag *NCGammaFix*. Thus, unlike backgrounds that are simulated using the MicroBooNE software suite as described in Sec. IIB, the NC coherent 1γ is simulated in a standalone GENIE simulation, and the outputs from GENIE simulation are fed into the MicroBooNE software suite for a full simulation of the final-state particle propagation and detector response. These events are then added to the nominal MicroBooNE neutrino event rate prediction used for this analysis. Figure 1 shows the distribution of the simulated outgoing photons as a function of the true photon energy and its true angle with respect to the neutrino beam direction, highlighting the forward nature of the outgoing photon from the coherent interaction.

The analysis was performed using a blinding procedure. It was developed primarily using MC overlay sam-




FIG. 1: The area-normalized 2D distribution of the true energy and angle with respect to the beam of outgoing photons in NC coherent 1γ events. This highlights the phase space of the outgoing photon, which populates the forward region and has a peak energy at ~ 0.3 GeV.

ples, while a 10% subset of the data was also accessible during the analysis development stage for validation purposes. A sideband sample was additionally defined (see Sec. VII A), and the analysis was fully frozen and validated on the sideband data sample before the signal region was unblinded. When the statistical significance of the observation was evaluated, two independent high-statistics π^0 measurements were used to constrain the predicted background rate and systematic uncertainties in the signal region. Section IV details the event selection, Sec. V describes the systematic uncertainties involved and how they are incorporated in this analysis, and Sec. VI shows the power of the π^0 constraint. The results obtained with the Run 1-3 data are discussed in Sec. VII A and VII B for the sideband and signal region, respectively.

IV. EVENT SELECTION

Driven by the rarity of the targeted signal, a series of selections, including traditional and gradient-boosted decision-tree (BDT) based selections, were developed with the goal of maximizing signal efficiency and minimizing background efficiency. The following sections present the detailed selections and tools developed for this analysis.

A. Event Breakdown

MC simulations are broken down into different categories based on truth information, in a manner similar to the analyses in Refs. [15] and [49], and shown in Fig. 2. For ease of reading, we reproduce the definitions of different categories here:

- NC coherent 1γ : This is the targeted signal. It includes all NC coherent single-photon production on argon inside the active TPC volume, regardless of incoming neutrino flavor.
- NC $1\pi^0$: All NC interactions in the active volume that produce one exiting π^0 regardless of incoming neutrino flavor. This category is further split into two sub-categories, “NC $1\pi^0$ Coherent” and “NC $1\pi^0$ Non-Coherent” contributions, based on their interaction types with coherent production making up 2.3% of all NC $1\pi^0$ interactions in GENIE. Non-Coherent scattering occurs when a neutrino interacts with a nucleon inside the argon nucleus, potentially knocking out one or more nucleons. NC $1\pi^0$ Coherent, like the NC coherent 1γ signal, is produced from coherent scattering of the neutrino with the nucleus, and the resulting π^0 tends to be very forward relative to the incoming neutrino beam. While NC $1\pi^0$ coherent is the closest in appearance to our signal, it is expected to be $\mathcal{O}(3\%)$ of all NC $1\pi^0$ events and the non-coherent contribution remains our largest background.
- NC $\Delta \rightarrow N\gamma$: NC single-photon production from the radiative decay of the $\Delta(1232)$ baryon in the active volume. This is the largest contribution of NC single-photon production relevant to the BNB beam. In this analysis, this is further broken down into two exclusive subsets, NC $\Delta \rightarrow N\gamma$ (1+p) and NC $\Delta \rightarrow N\gamma$ (0p), depending on the outgoing nucleon and its kinetic energy (KE). NC $\Delta \rightarrow N\gamma$ (1+p) refers to NC Δ radiative decay events that have at least one proton exiting the nucleus with KE larger than 50 MeV. $\Delta \rightarrow N\gamma$ (0p) is the complementary subset and includes NC Δ radiative decay events with either no protons exiting the nucleus, or protons exiting the nucleus but with KE < 50 MeV. In all cases neutrons of any energy are allowed to exit the nucleus.
- CC $\nu_\mu 1\pi^0$: All ν_μ CC interactions in the active volume that have exactly one true exiting π^0 . CC $\bar{\nu}_\mu 1\pi^0$ is not included in this category but in “BNB Other,” defined below.
- CC $\nu_e/\bar{\nu}_e$ Intrinsic: All CC ν_e or $\bar{\nu}_e$ interactions in the active volume.
- BNB Other: All remaining BNB neutrino interactions that take place in the active TPC volume and are not covered by the above five categories, such as multiple π^0 events and η meson decay.
- Dirt (Outside TPC): All BNB neutrino interactions that take place outside the MicroBooNE active TPC but have final states that enter and deposit energy inside the active TPC detector. This can originate from scattering off liquid argon in the cryostat vessel outside the active TPC volume or

from interactions in the concrete and “dirt” surrounding the cryostat itself.


- Cosmic Data: Coincident cosmic ray interactions that take place during a BNB beam-spill but without any true neutrino interaction present.

B. Topological Selection and Preselection


The event selection starts with a topological selection on outputs from Pandora, specifically requiring one reconstructed shower and zero reconstructed tracks ($1\gamma 0p$), which is the expected topology for NC coherent 1γ signal in the detector. This topological requirement selects 28.1% of simulated NC coherent 1γ events; the remainder of the events do not pass the selection cut because 32.7% of the true signal events are reconstructed with no track or no shower by Pandora (i.e no candidate neutrino interaction), while 39.2% are reconstructed with either ≥ 1 tracks or ≥ 2 showers. Then, preselection cuts are applied in order to remove obvious backgrounds. The preselection cuts require that the reconstructed shower energy be larger than 50 MeV to remove Michel electrons from muon decays. Additionally a fiducial volume cut of at least 2 cm away from the space charge boundary [51] is made on the reconstructed shower start point, in order to reduce the number of partially contained showers. Relative to the topological stage, the preselection cuts further remove 33.9% of the overall background—mostly dirt events and cosmic ray background—while preserving 98.6% of the signal. At the preselection stage, 27.7% of the NC coherent 1γ signal remains out of all simulated signal events in the TPC, and the corresponding signal to background ratio is about 1 : 2700. Figure 2 shows the predicted NC coherent 1γ signal scaled by a factor of 2500 and the nominal predicted background at preselection, highlighting the rarity of the signal and the different regions of phase space in reconstructed shower energy populated by the signal and backgrounds.

C. Event-level Boosted Decision Tree Based Selection

At preselection, BDTs are developed and optimized to further differentiate the signal from background events. There are in total six (6) tailored BDTs developed using the XGBoost library [52], each trained with variables derived from reconstructed objects. Every BDT targets a different background and thus is trained with a different category of background events. There are four BDTs trained with event-level variables, which are variables characterizing the event as a whole, and two BDTs trained with cluster-level variables associated with clusters of reconstructed hits in each event. Simulated MC events are separated into statistically exclusive sets; test sets independent from the BDT training sets are used to



(a)



(b)

FIG. 2: Predicted background distributions and observed data as a function of (a) the reconstructed shower energy and (b) the reconstructed shower angle with respect to the neutrino beam, at preselection. At this early stage many categories contain too few events to observe, including the coherent signal. As such, the orange histogram overlaid (not stacked) shows the distribution of NC coherent 1γ signal scaled by 2500 in (a) and 1000 in (b) for it to be visible. At this stage, the signal and backgrounds populate different regions of phase space in reconstructed shower energy, with the ~ 300 MeV peak in signal being dwarfed by low energy shower activity. In the angular phase space one can clearly see the forward nature of the signal compared to the downward nature of cosmic-originating showers.

obtain the MC prediction, which later is used to choose the cut values and perform data-MC comparisons. In this section we discuss the event-level BDTs; the formation of clusters and cluster-level BDTs are described in Sec. IV D.

The largest background contributions remaining at preselection are cosmic-ray background, dirt events ($\sim 50\%$ of which are neutrino interactions containing a fi-

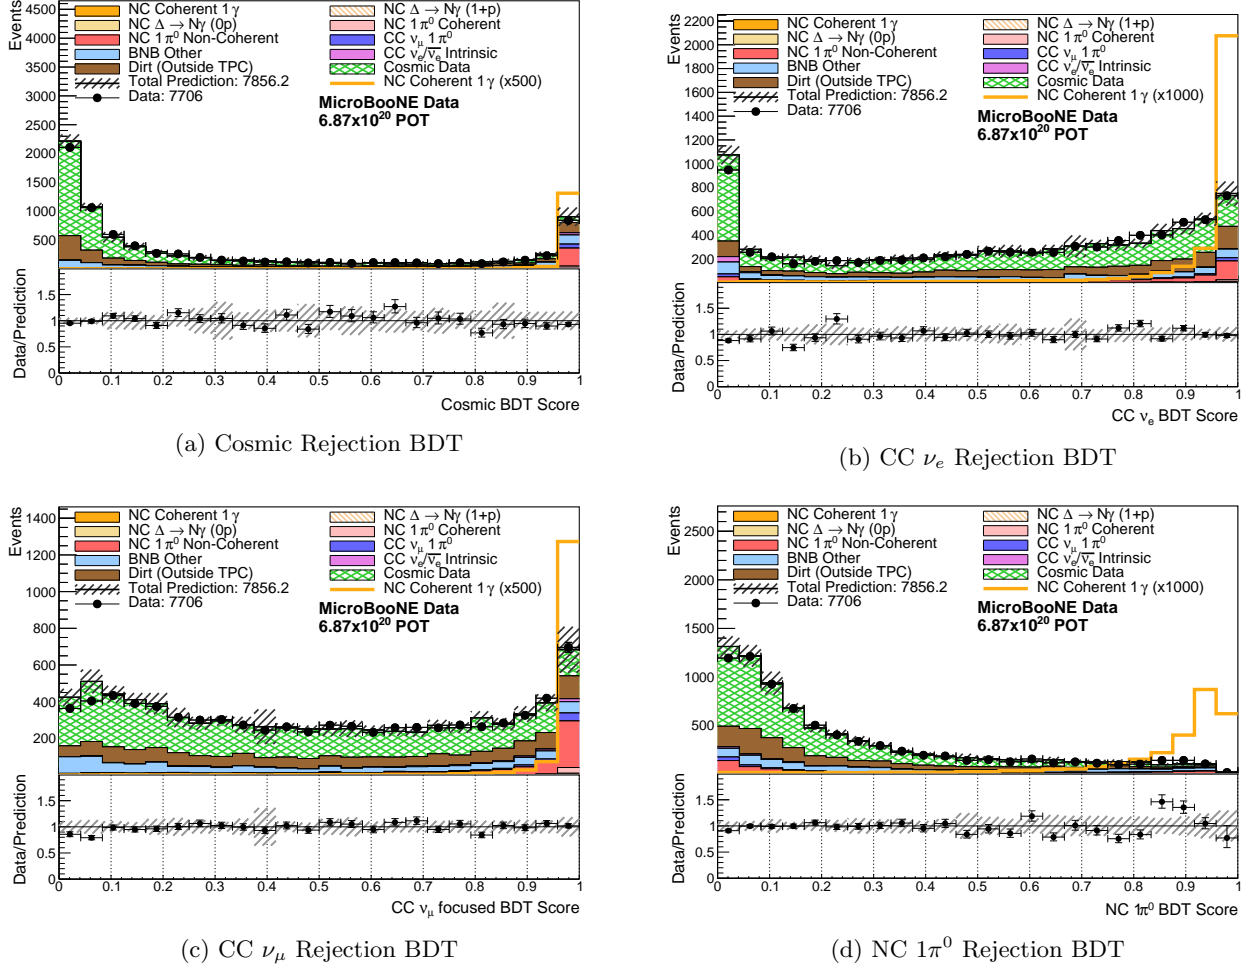


FIG. 3: The BDT classifier scores for (a) the Cosmic BDT; (b) CC ν_e BDT; (c) CC ν_μ focused BDT; and (d) NC $1\pi^0$ BDT. Higher scores indicate more NC coherent 1γ signal-like events. The stacked histograms show the predicted background distribution, while the orange histogram (not stacked) represents the distribution for the NC coherent 1γ signal, both at preselection stage. The signal histogram is scaled for visual purposes by a factor of 1000 in (c) and (d) and 500 in (a) and (b).

nal state π^0), BNB Other, and NC non-coherent $1\pi^0$, listed in decreasing order of the size of the predicted contribution. Three event-level BDTs are employed to target these backgrounds, and another BDT is designed to specifically target the electron background from CC interactions induced by the intrinsic $\nu_e/\bar{\nu}_e$ in the beam. To enhance the ability of BDTs to learn the differences between NC coherent 1γ events and backgrounds, all four BDTs are trained with selections of events passing the preselection requirements. The same NC coherent 1γ overlay sample is used as the “signal” during training of the four BDTs with, in addition to the preselection, the requirements that the reconstructed shower originates from photon activity in truth and $> 50\%$ of the reconstructed hits in the shower result from energy deposition of the simulated activity. The background definitions used and the key features in training the four BDTs are described below:

a. Cosmic BDT: The goal of the cosmic BDT is to differentiate the NC coherent 1γ signal from misidentified cosmic backgrounds. The cosmic BDT trains on BNB external data as background, and makes use of the fact that cosmic rays usually travel from the top to the bottom of the detector, resulting in reconstructed showers that are oriented more vertically and have less extent in the beam direction when compared to that of the signal.

b. CC ν_e BDT: The CC ν_e BDT aims to remove reconstructed showers originating from electrons instead of photons, and trains on simulated, preselected CC $\nu_e/\bar{\nu}_e$ events. One key handle for photon/electron separation is the shower dE/dx . Most photons in the energy range relevant to this analysis lose energy through e^+e^- pair-production; these then further radiate, producing an electromagnetic shower. For electron showers, dE/dx at the shower start is similar to that of a minimum ionizing particle (MIP), around 2 MeV/cm, whereas the photon

shower start typically has a $dE/dx \sim 4$ MeV/cm.

c. CC ν_μ focused BDT: The CC ν_μ focused BDT aims to remove any backgrounds other than the cosmic, NC Δ radiative decay, CC $\nu_e/\bar{\nu}_e$ and NC $1\pi^0$. Among these background events, about 76% are CC $\nu_\mu/\bar{\nu}_\mu$ events, 60% of which are CC events without a π^0 exiting the nucleus, and for which the muon is misreconstructed as a shower. Since muons are minimum ionizing, variables such as shower dE/dx , average energy per hit in the reconstructed shower, and the Pandora shower score are important for the separation of misidentified muons from true coherent single photons. The CC ν_μ focused BDT is trained with simulated $\nu/\bar{\nu}$ -interactions other than NC Δ radiative decay, CC $\nu_e/\bar{\nu}_e$ and NC $1\pi^0$ as background events.

d. NC $1\pi^0$ BDT: The NC $1\pi^0$ background is harder to remove compared to cosmic and other ν -induced backgrounds because in most cases the reconstructed shower in misidentified NC $1\pi^0$ events is indeed from a true photon (most likely the leading photon from π^0 decay). There are a variety of reasons that can lead to single shower reconstruction in the π^0 sample: 1) The second shower from π^0 decay is not visible in the detector; this could happen when the photon is absorbed by the medium, leaves the detector before pair-producing, or is not of sufficient energy to be detected. 2) The second shower deposits energy in the detector but the 3D reconstruction fails; this could happen when the cluster-matching across planes in Pandora fails. 3) The hits from the second photon are reconstructed in 3D but not associated with the neutrino interaction; this happens when cosmic-rays interfere with the energy deposition on one or more planes. To mitigate failure case (2), a cluster-level second-shower veto (SSV) BDT is developed following the exact same approach as in Ref. [15] to identify the possible presence of the second shower that is not reconstructed. The resulting SSV BDT score is not used as an independent variable to directly decide whether to retain or reject an event. Instead, it serves as a training variable input for the NC $1\pi^0$ BDT to improve its π^0 rejection efficiency. The SSV BDT is discussed in more detail in the next section, and a selection of the SSV BDT outputs are used to train the NC $1\pi^0$ BDT. In addition to variables involving the missing second shower, variables associated with the primary reconstructed shower that contribute the most in separating the signal from NC $1\pi^0$ backgrounds are the reconstructed shower energy, angle, and projected momentum on the vertical plane perpendicular to the neutrino beam direction.

The outputs of the event-level BDTs are scores in the range of [0,1]; a higher event-level BDT score indicates the event is more NC coherent 1γ signal-like. Figure 3 shows the predicted BDT score distributions at preselection. The NC coherent 1γ signal is scaled up in order to highlight the separation between the signal and the targeted background in the BDT responses.

Selections for these four event-level BDTs are chosen by placing cuts on each resulting BDT score in order to

maximize the statistical significance of the NC coherent 1γ signal, i.e. $\frac{N_{\text{sig}}}{\sqrt{N_{\text{bkg}}}}$, where N_{sig} is the number of signal events and N_{bkg} is the number of background events. The possible selection cut values are explored simultaneously on all four BDTs and the signal significance is evaluated for every choice of BDT cut values. Significant degeneracy is found in the four-dimensional space of possible selection cut values. The selection cuts chosen to maximize signal selection efficiency within the degenerate region are shown in Tab. I. The set of events for which their BDT scores are higher than the corresponding selection values constitutes the ‘‘Single-Photon’’ rich selection.

Event-level BDTs	Optimized BDT Score Value
Cosmic	0.990
CC ν_e	0.885
CC ν_μ	0.992
NC $1\pi^0$	0.891

TABLE I: Single-Photon selection values on the four primary event-level BDTs, optimized based on the signal statistical significance ($\frac{N_{\text{sig}}}{\sqrt{N_{\text{bkg}}}}$).

D. Cluster-level Boosted Decision Tree Based Selection

Cluster-level BDTs are designed to identify activity of interest that is missed by Pandora pattern recognition. There are two cluster-level BDTs employed in this analysis. First, a second-shower veto BDT targets events that are missing the second shower from π^0 decay. Second, a proton-stub veto (PSV) BDT aims to identify proton activity near the shower vertex, in order to remove non-coherent backgrounds. Both BDTs build on hits that are not reconstructed in 3D by Pandora; individual scattered hits are clustered together to form candidate clusters of interest which then get classified by the BDTs. We will discuss the clustering mechanism first before describing the specifics of each BDT.

1. Hit Clustering

Each reconstructed hit consists of a calibrated charge integral and associated wire and timing information. Nearby neighboring hits are grouped together on a plane-by-plane basis; this is referred to as ‘‘hit clustering’’. Hits with lower than 25 charge integral units are considered ‘‘noise’’ hits and ignored during clustering formation. The separation of two hits on the same plane is defined by the following metric:

$$D = \sqrt{((w_1 - w_2) * 0.3 \text{ cm})^2 + (\frac{t_1 - t_2}{25 \text{ cm}^{-1}})^2}, \quad (2)$$

where w_1, w_2 are the wire number of two hits, t_1, t_2 are the associated time tick (1 tick = 0.5 μs) information and

where the normalization factors are chosen to make the wire spacing and time interval difference comparable in physical extent.

The clustering is performed using the density-based spatial clustering of applications with noise (DBSCAN) algorithm [53] with two required parameters:

- N_{minPts} : the minimum number of hits required to form a cluster.
- ϵ : the maximum distance between two hits. This distance requirement is used in both initial cluster formation, which requires at least N_{minPts} hits all within a distance of ϵ from each other, and cluster expansion, where an additional hit is absorbed as long as it is within a distance of ϵ from at least one hit in an existing cluster.

The result of DBSCAN is a group of clusters on three planes that are robust to spatially outlying hits and arbitrarily shaped, complying with the restriction from ϵ .

2. Second Shower Veto BDT

The SSV BDT, as mentioned in Sec. IV C, targets the second shower from π^0 decay in the detector, which has an expected topology of a sparse cascade of hits near the primary reconstructed shower. Thus, the candidate clusters for the second shower are formed with $N_{\text{minPts}} = 8$ and $\epsilon = 4$ cm. Undesirable clusters can result from electronic noise along a single wire during continuous time intervals and from correlated noise across several wires at the same time. To avoid clustering on such noise, the candidate clusters are required to span at least four time ticks and three different wires. Further, the spatial variance of hits within a cluster along any single axis cannot exceed 99.9% of the total spatial variance of the cluster. This is evaluated by performing principal component analysis (PCA). This further reduces the misclustering of noise and yields one or more second shower candidate clusters per plane per event.

The reconstructed calorimetric and spatial variables associated with second shower candidate clusters are the inputs to the SSV BDT. More specifically, second shower candidate clusters that are truth-matched to a different photon from the primary reconstructed shower in the NC $1\pi^0$ sample serve as the training signal, while second shower candidate clusters formed in simulated NC coherent 1γ events are used as training background. The output of the BDT is a score assigned to each cluster from 0 to 1. The higher the score, the higher the probability that the cluster originates from a photon from π^0 decay. The output of the SSV BDT is not itself directly cut on in event selection; instead, the maximum SSV BDT score of all individual clusters on each of the three planes (i.e. the most likely second-shower candidate per plane) is used as input variables to the NC $1\pi^0$ BDT. Figure 4 shows the distribution of the maximum SSV BDT scores of all clusters formed in events.




FIG. 4: Distribution of the maximum SSV BDT score of all clusters formed in events at preselection. Nominal predictions for backgrounds are stacked in the colored histograms, while the signal prediction is scaled by a factor of 1000 and plotted separately in the orange histogram (not stacked). A higher score means a higher probability of a second shower being present in the event. Note that, if no second shower candidate cluster is formed in an event, the maximum SSV BDT score for this event is set to zero.

3. Proton Stub Veto BDT

Besides the SSV score, another handle is leveraged to aid in the rejection of background: small traces of protons. Pandora can consistently perform particle identification on proton tracks starting at around 40 MeV of kinetic energy, however, hand scanning of NC Δ radiative decay events shows visible true low momentum proton-like activity on one or more planes near the reconstructed shower in some $1\gamma 0p$ events. This proton-like activity usually sits very close to the backward projection of the shower direction. The PSV BDT is designed to identify such indications of protons missed by reconstruction, through which we can remove non-coherent backgrounds, specifically misidentified NC backgrounds with protons exiting the nucleus.

The building block of the PSV BDT is proton candidate clusters, similarly to the SSV BDT. Since most protons that are missed by the reconstruction have low energy, and proton tracks are expected to be very straight with dense energy deposition, proton candidate clusters are defined by requiring that they are dense, specifically by setting $\epsilon = 1$ cm during the DBSCAN clustering. Driven by handscan observations, we also set N_{minPts} to 1 to allow the capture of tiny clusters originating from proton activities. This allows small noise hits to leak in, leading to $\mathcal{O}(100)$ proton candidate 2D clusters per event, approximately evenly distributed per plane. Thus the PSV BDT must be efficient in correctly identifying true proton clusters from a substantial amount of background clusters arising from other activity. Two handles help

solve this problem: the geometric relation between true proton candidate clusters and the reconstructed shower, and the distinct calorimetric profile of a proton track.

The cartoon in Fig. 5 shows a potential geometric relation between a proton candidate cluster and the reconstructed shower on any given plane. In backgrounds that have true protons exiting the nucleus and a reconstructed shower from a true photon, the proton cluster is expected to intersect the line of backward projection of the reconstructed shower. This means the minimum perpendicular distance from the candidate cluster to the direction of the primary reconstructed shower, i.e. the impact parameter of the cluster, should be zero for the proton candidate clusters. Additionally, the PSV BDT takes advantage of the straightness and the Bragg peak signature expected for a proton track to identify true proton clusters.

A sample of simulated NC Δ radiative decay events is used to train the PSV BDT: proton candidate clusters in the NC Δ sample that are truly from protons are used as training signals, while all remaining candidate clusters, be it from noise, cosmogenic origin or non-proton neutrino induced activity, are the training background. Unlike the SSV BDT, which is not used as a selection cut, a cut is applied directly on a derived PSV BDT output to further improve the signal-to-background ratio after the single-photon selection.

The derived PSV BDT output used is the maximum PSV BDT score of all proton candidate clusters on plane 0 and plane 2 (abbreviated as “maximum PSV score on planes 0 and 2”). The reason for not including clusters formed on plane 1 is due to the fact that an excess amount of proton candidate clusters is found in data in the open dataset. The excess is localized in specific regions of plane 1 and found to be due to a mismatch of run periods used in the MC overlay sample and the open data. This was resolved by requiring the overlay sample in simulation and the collected data to come from overlapping run periods. However, to be conservative, the PSV BDT score information on plane 1 is not used.

Figure 6 shows the distribution of the maximum PSV score on planes 0 and 2 for the predicted background at preselection. As expected, the NC coherent 1γ signal and most of the background pile up on the left, while the right corner is most populated by the NC non-coherent $1\pi^0$. The distribution of the same variable with the cosmic, CC ν_e and CC ν_μ -focused BDT requirements applied is shown in Fig. 7. The resulting sample is a photon-rich sample dominated by the NC non-coherent $1\pi^0$ s. The peak near 1 is mostly populated by NC non-coherent $1\pi^0$ background, highlighting the strong separation power of the PSV BDT. The NC coherent 1γ signal and majority of the NC coherent $1\pi^0$ background cluster at low BDT score regions due to their coherent nature, as expected.

The requirement on the maximum PSV score on planes 0 and 2 is applied to events in the single-photon selection, and the cut value is optimized to maximize the signal significance with the neutrino flux and neutrino interaction uncertainties (discussed in Sec. V) included and




FIG. 5: *Top*: Cartoon showing the relative position of a proton candidate cluster to the primary reconstructed shower on a plane. Two important variables for the PSV BDT are also highlighted: the minimum distance between the proton cluster and the reconstructed shower start point, and the impact parameter of the candidate cluster to the back-projection of the reconstructed shower direction. If the proton stub candidate is indeed from a proton that comes from the same vertex as the reconstructed shower, the impact parameter is expected to be close to zero, assuming the reconstructed shower direction is accurate. *Bottom*: An example data event display showing a clear candidate proton stub in the back-projection of the reconstructed shower on plane 2, with a high PSV BDT score of 0.94 and a reconstructed energy of 13.5 MeV.

constrained from the 2γ samples (discussed in Sec. VI). This optimized selection value is found to be 0.2. Thus, events that have a maximum PSV score on planes 0 and 2 less than 0.2 comprise the the “Coherent-Rich Subset” of the single-photon selection.

Figure 8 shows the rejection efficiency of the PSV BDT on simulated NC non-coherent $1\pi^0$, NC $\Delta \rightarrow N\gamma$ ($1+p$), and NC $\Delta \rightarrow N\gamma$ ($0p$) events. The rejection efficiency




FIG. 6: The maximum PSV score among all clusters on plane 0 and plane 2 for events at preselection. Higher score indicates increasing confidence of a proton exiting the nucleus that is missed by Pandora. Nominal predictions for backgrounds are stacked in the colored histograms, while the signal prediction is scaled by a factor of 1000 and plotted separately in the orange histogram (not stacked).




FIG. 7: The maximum PSV score among all clusters on plane 0 and plane 2 for events passing the three event-level BDTs (cosmic, CC ν_e and CC ν_μ focused BDTs). While there can be many clusters per event, this plot is showing just the maximum so there is one entry per whole event. The nominal prediction for the signal is shown on top of the background stacked histograms as an orange filled histogram, which can be seen in the first bin on the left.

is evaluated by calculating the fraction of events that fail the PSV BDT requirement in each bin, relative to selected events passing the topological, preselection, and cosmic, CC ν_e and CC ν_μ BDT requirements. We observe that the PSV yields a high rejection efficiency on protons at low energies on all three samples, with an average of approximately 70% rejection efficiency for protons with





FIG. 8: The rejection efficiency of the PSV BDT as a function of the KE of the leading proton exiting the nucleus for both non-coherent NC π^0 as well as NC $\Delta \rightarrow N\gamma$ events. The error bar represent the uncertainty arising from finite MC statistics. The vertical line at 50 MeV KE is the nominal threshold for Pandora to reconstruct a proton track for comparison.

KE less than 50 MeV. Due to the fact that each event has many candidate proton clusters coming from noise, cosmic induced activity, as well as EM-induced clusters associated with the reconstructed photon shower itself, a certain number of events will be vetoed due to random coincidence. To estimate this frequency we study our true coherent signal, which contains no protons, and note that only 14.8% of these events are vetoed, indicating a dominant portion of the $\approx 70\%$ veto efficiency seen in Fig. 8 is truly due to observing protons and not due to clusters in random coincidence.



E. Coherent-Rich Subset

We first show the results at the single-photon selection stage in three variables that highlight the reconstructed shower energy and angle, as well as the PSV features of signal and backgrounds in Fig. 9. The selection has a reconstructed shower distribution with an energy peak of around 300 MeV and a very forward angle, which is expected for the NC coherent 1γ signals. As expected, NC non-coherent $1\pi^0$ events are the most dominant background, and the PSV distribution at this stage motivates the additional requirement on the PSV variable to be less than 0.2.

Table II shows all background and signal predictions for the single-photon selection and the coherent-rich subset selection, normalized to 6.87×10^{20} POT. Prior to the PSV cut there are a total of 54.5 predicted background events and 1.3 signal events, leading to a signal-to-background ratio of 1:42. The PSV requirement further improves the signal-to-background ratio to 1:30 by filtering out almost half of the NC non-coherent $1\pi^0$ back-



(a) Reconstructed Shower Energy

(b) Reconstructed Shower $\cos(\theta)$ 

(c) Maximum PSV score between Plane 0 and Plane 2

FIG. 9: Distributions of the single-photon selection in: (a) reconstructed shower energy; (b) reconstructed shower $\cos(\theta)$; and (c) the maximum PSV score on planes 0 and 2. The selected events have a reconstructed shower with an energy peak at ~ 300 MeV in a very forward direction, which is expected for the NC coherent 1γ signal.

ground. Figure 10 shows the selection efficiency for the simulated NC coherent 1γ signal in the active TPC at different stages of the analysis, as a function of the true energy and angle of the outgoing photon.

Stage	Single-Photon Selection	Coherent-Rich Subset
NC coherent 1γ (Signal)	1.3	1.1
NC $\Delta \rightarrow N\gamma$ (1+p)	0.3	0.1
NC $\Delta \rightarrow N\gamma$ (0p)	2.5	1.3
NC $1\pi^0$ Non-Coherent	29.9	15.8
NC $1\pi^0$ Coherent	3.8	3.1
CC ν_μ $1\pi^0$	5.0	3.6
CC ν_e and $\bar{\nu}_e$	0.6	0.4
BNB Other	3.5	2.9
Dirt (outside TPC)	6.4	4.8
Cosmic Ray Data	2.4	0.8
Total Prediction (Unconstr.)	55.8	34.0
Total Prediction (Constr.)	45.8	29.0

TABLE II: The expected event rates in the single-photon and coherent-rich subset selections for all event types, normalized to the Run 1-3 data POT (6.87×10^{20}). The total, constrained prediction is evaluated with the conditional constraint from the $2\gamma 1p$ and $2\gamma 0p$ samples.

The selection efficiencies for all signal and background categories at different analysis stages are included in Table III, calculated relative to predictions after the topological selection. We note again that this initial topological efficiency in which a single shower is reconstructed is 28% for our signal. The event-level BDTs successfully reject all types of backgrounds by a large fraction, especially the intrinsic CC ν_e , BNB Other, dirt and cosmic backgrounds. Transitioning from the single-photon selection stage to the coherent-rich subset, the cut on the PSV BDT acts mostly on the NC $\Delta \rightarrow N\gamma$ (1+p), NC $\Delta \rightarrow N\gamma$ (0p), and NC non-coherent $1\pi^0$ backgrounds, with corresponding rejection efficiencies of 68.8%, 47.9%, and 47.4%, respectively. Note that NC $\Delta \rightarrow N\gamma$ (0p) consists of NC Δ radiative decay events with neutrons exiting the nucleus as well as NC Δ radiative decay events with exiting protons with KE < 50 MeV, the latter of which is specifically targeted by the PSV. Overall, relative to the topological selection, the coherent-rich subset selection achieves a 97.5% rejection efficiency on NC non-coherent $1\pi^0$ and a 99.99% rejection efficiency on the cosmic background while keeping the signal efficiency at 39%.

V. SYSTEMATIC UNCERTAINTIES

The sources of systematic uncertainty considered in this analysis are broken down into contributions from 1) neutrino flux, 2) neutrino interaction cross sections, 3) hadron-argon reinteractions, 4) detector response, and 5) finite MC statistics.

Category	Preselection Eff. [%]	Single-Photon Eff. [%]	Coherent-Rich Subset Eff. [%]
NC Coherent 1γ	98.59	45.85	39.09%
NC $\Delta \rightarrow N\gamma$ (0p)	98.30	19.18	9.99
NC $\Delta \rightarrow N\gamma$ (1+p)	97.13	7.99	2.49
NC $1 \pi^0$ Coherent	96.91	8.38	6.83
NC $1 \pi^0$ Non-Coherent	96.39	4.75	2.50
CC $\nu_\mu 1\pi^0$	90.21	3.07	2.24
BNB Other	82.17	0.30	0.25
CC $\nu_e/\bar{\nu}_e$ Intrinsic	94.59	0.69	0.47
Dirt (Outside TPC)	61.67	0.25	0.18
Cosmic Data	61.01	0.03	0.01

TABLE III: Selection efficiencies for our simulated signal and background categories at different stages of the analysis, calculated with respect to the number of selected events after the topological requirements.

The flux uncertainty includes proton delivery, hadron production, hadronic interactions, and the modeling of the horn magnetic field [31]. Both the intensity and position of the proton beam when it reaches the BNB target is monitored, and the uncertainty on our expected neutrino rate due to proton delivery is estimated to be a 2% normalization uncertainty. The hadronic production concerns the production of secondary particles (π^\pm , K^\pm , K_L^0) which decay to produce neutrinos. Hadronic interaction cross sections impact the hadron production rate in the primary proton-Be interaction, and the rate of pion absorption in the target and the horn, which ultimately affects the rate and shape of the neutrino flux. Uncertainty on the horn magnetic field modeling comes from two sources: the intrinsic variation in the horn current, and uncertainty in the modeling of the current within the inner cylinder in the horn. Tools and techniques developed by MiniBooNE [31] are adopted to evaluate flux uncertainty in MicroBooNE. The “multisim” technique is employed for flux uncertainty evaluation where multiple MC simulated universes are generated and, in each universe, events are reweighted independently based on neutrino parentage, neutrino type, and energy.

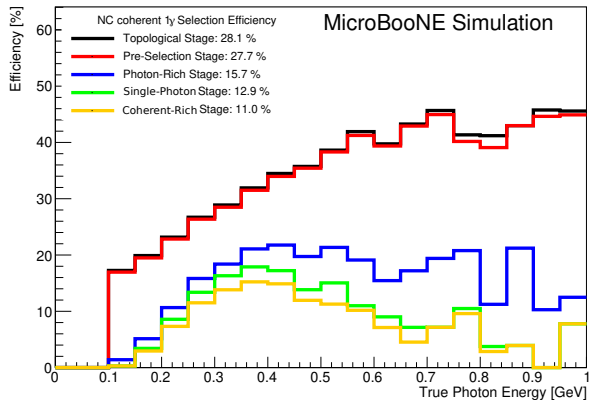
The neutrino interaction cross section uncertainties arise from uncertainties in the model parameters used in the GENIE prediction. This includes uncertainty in NC resonant production, the branching ratio, and the angular uncertainty on NC Δ radiative decay, uncertainties on CC quasielastic and resonance production, as well as other NC and CC interactions, and final state interactions. More details on the treatment of cross section parameter uncertainties can be found in Ref. [33]. Simulated events are reweighted using GENIE reweighting tools when model parameters are varied. For the cross section uncertainty evaluation, there exist universes resulting from single parameter variations for certain model parameters as well as universes with multiple parameters varied simultaneously, which allows incorporating correlations between model parameters. Due to technical limitations, the NC and CC coherent π production cross sections are assigned with a $\pm 100\%$ normalization uncertainty.

Hadron-argon reinteractions refer to the scattering of

charged hadrons off the argon nuclei through hadronic interactions during propagation, which could significantly bend the particle trajectory or lead to additional particle production. GEANT4 is used to simulate hadrons propagating through the detector medium [54], and events are reweighted independently for π^+ , π^- , and proton interactions via the GEANT4REWEIGHT tool [55].

The detector systematic uncertainty stems from the difference between detector simulation and the actual detector response. There are three main types of uncertainties: 1) uncertainty in TPC waveform from the electronic response simulation, 2) uncertainty in the light simulation, and 3) other detector effects modifying the distribution of drifting ionization electrons. MicroBooNE employs a novel data-driven technique to evaluate the uncertainty on TPC wire waveforms, by parameterizing the differences between simulated cosmic and observed cosmic data at the level of Gaussian hits after waveform deconvolution [56]. This avoids the large computation required for detector simulation and signal deconvolution and allows the detector uncertainty to be evaluated in a model-agnostic way. Four variations as functions of X position (drift direction), YZ position (vertical and beam direction), and the angular variables θ_{XZ} and θ_{YZ} for the particle trajectory, respectively, are considered. Variations associated with the light yield consider an overall 25% reduction in the light yield, variation in the light attenuation, and changes in Rayleigh scattering length. Other detector effects, including the space charge effect due to the buildup of argon ions [44, 57], and variations in the ion recombination model [58] are accounted for through additional variations. Independent central value (CV) MC simulations, apart from the samples used in training the BDTs and making the final predictions, are generated together with the corresponding detector variations discussed above. The detector CV and variation samples are run through the reconstruction and analysis selection chain before being compared for detector uncertainty evaluation.

The systematic uncertainties are encapsulated in covariance matrices, which describe the covariance across bins in a predicted spectrum or across different predicted spectra. For each systematic error and associated varia-



(a) Selection Efficiency v.s. True Photon Energy


(b) Selection Efficiency v.s. True Photon $\cos \theta$

FIG. 10: Efficiencies at various stages of the selection for the NC coherent 1γ signal as a function of (a) true photon energy and (b) true photon angle with respect to the neutrino beam, highlighting the photon phase space targeted by this analysis. The “photon-rich stage” refers to the stage where the three event-level BDTs (cosmic, CC ν_e and CC ν_μ focused BDTs) are applied. The efficiencies are calculated as the ratio of the selected NC coherent 1γ events within the range plotted for each variable to the total prediction of NC coherent 1γ .

tion(s), a covariance matrix is constructed via:

$$M_{ij} = \frac{1}{N} \sum_{k=1}^N (n_i^k - n_i^{CV})(n_j^k - n_j^{CV}) \quad (3)$$

where M_{ij} is the resulting covariance matrix describing the covariance between bin i and bin j in the predicted spectrum, N is the number of universes generated for a given systematic error source, n_i^k is the prediction in bin i in the k th universe and n_i^{CV} is the prediction in the same bin of the CV simulation. For reweightable systematics, the n^k spectrum is calculated by reweighting

simulated events in the selection, and for the detector systematics, the n^k spectrum is the spectrum of the variation sample after applying all selection cuts discussed in Sec. IV. The breakdown of uncertainty contributions into the four main categories for the single-photon selection is shown in Fig. 11 as a function of the reconstructed shower energy. The dominant uncertainty contribution at the single-photon selection stage is the uncertainty in detector modeling, followed by the uncertainty in cross section modeling. As this is a rare signal search, the MC statistics in the detector variation samples can get quite low in a given bin when one considers a binned spectra as in Fig. 11. Thus, for the purpose of a final fit to search for coherent single-photon events we use a single bin, integrating over the entire shower energy range. The result of this can be found in Fig. 12, where MC statistics in the detector variation samples are less of a concern, resulting in an overall lower detector systematic uncertainty. Nonetheless, cross section and detector uncertainty remain the two largest contributions.

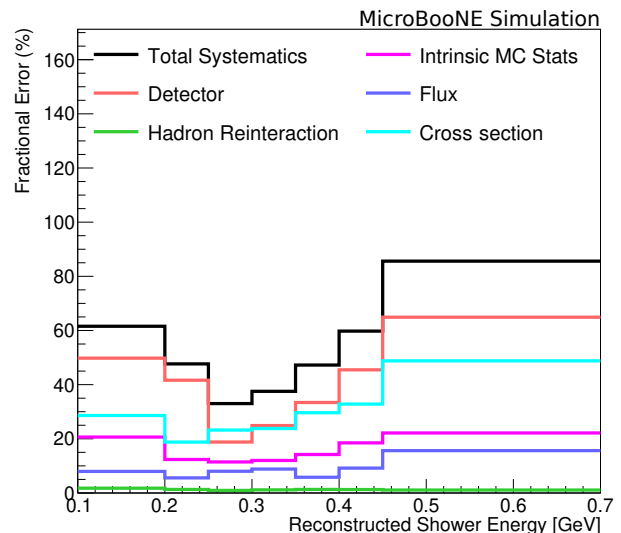


FIG. 11: Overall systematic uncertainty at the single-photon selection stage, as a function of the reconstructed shower energy. The contribution from each source is highlighted in different colors and the quadrature sum of all systematics is shown in black.

VI. π^0 CONSTRAINT

The goal of this analysis is to search for the NC coherent 1γ process in the observed data. This is quantified by a normalization scaling factor x of the nominal SM predicted rate for this process. Extraction of the parameter x and its range is achieved by fitting the MC prediction to the observed data while allowing the scaling factor x to vary. In order to yield higher sensitivity to the parameter





FIG. 12: Overall systematic uncertainty for the coherent-rich subset selection, with contribution from each source highlighted in different colors, and in black is the sum in quadrature of all systematics. Driven by the low statistics available at the coherent-rich stage, a single bin is used.


x , external high-statistics measurements are used in the fit to constrain the predicted background and systematic uncertainty in the signal region, in a manner similar to MicroBooNE’s LEE analyses [15, 59–62].

Given its dominant contribution, NC $1\pi^0$ selections, developed for MicroBooNE’s Pandora-based NC Δ radiative decay search (detailed in Ref. [49]), are chosen as the constraining sample for this analysis, with $\approx 20\%$ more data than was used before in Ref. [49]. There are two $1\pi^0$ selections: $2\gamma 1p$ and $2\gamma 0p$. While both of these selections aim to reconstruct NC $1\pi^0$ events, $2\gamma 1p$ specifically targets π^0 events originating from resonant $\Delta \rightarrow N\pi^0$ decays while the $2\gamma 0p$ selection requires no reconstructed tracks.

We fit three samples (signal selection, $2\gamma 1p$ and $2\gamma 0p$ selection) simultaneously to maximize the sensitivity to the NC coherent 1γ process. The distributions of the 2γ samples used in the fit are shown in Fig. 13. Following the conditional constraint approach described in [60], we observe that the 2γ constraint reduces the MC background prediction from $34.0 \pm 8.4(\text{syst}) \pm 5.8(\text{stats})$ to $29.0 \pm 5.2(\text{syst}) \pm 5.4(\text{stats})$ for the coherent-rich signal selection, which corresponds to a $\sim 25\%$ reduction in the systematic uncertainty.



(a) $2\gamma 1p$



(b) $2\gamma 0p$


FIG. 13: Distributions of the $2\gamma 1p$ and $2\gamma 0p$ constraining samples as a function of the reconstructed π^0 momentum, shown in the binning used in the simultaneous fit. A mild deficit in data is observed in the $2\gamma 1p$ sample, but in general, good data-MC agreement is seen within the systematic uncertainty.

VII. RESULTS



A. Sideband

The analysis is designed as a blind analysis, where 10% of the total data during the first run period is open and utilized to develop the analysis and validate the agreement between data and MC in variables of interest and input variables to all BDTs involved. The analysis selections presented and the backgrounds involved are similar to those of Ref. [15], and the NC π^0 and BNB Other backgrounds with a single-shower topology are validated in Ref. [15]. The performance of the shower energy reconstruction and validation of NC π^0 modeling are demonstrated in the $2\gamma 1p$ and $2\gamma 0p$ NC π^0 measurements with the full dataset [49].

To validate the performance of the PSV BDT, after topological and preselection cuts, we also isolate a sideband using the full dataset that is proton-rich by reversing the cut on the NC π^0 BDT (<0.891) while relaxing cuts on other event-level BDTs: Cosmic BDT >0.8 , CC ν_μ BDT >0.8 , CC ν_e BDT >0.5 . A total of 850.0 events are predicted, of which 54.0% of non-cosmic MC events are predicted to have protons exiting the nucleus (regardless of proton KE). Consistency between data and MC is inspected through the goodness-of-fit test that incorporates systematic uncertainties and the Combined-Neyman-Pearson statistical uncertainty [63]. Figure 14 shows the distributions of the reconstructed shower energy, angle, and the maximum PSV score on planes 0 and 2 for the sideband sample. A mild deficit in data is observed in the region of high PSV score. The subsample of events with high PSV score (above 0.8) was isolated and found to be uniformly distributed in shower kinematic variables. Overall, the data-MC agreement for the PSV variable is good, with a corresponding p -value of 0.78 when all systematic uncertainties are taken into account. Event-level and cluster-level variables including all BDT input variables were also inspected, and found to give good agreement within systematic uncertainty.



(a) Reconstructed Shower Energy

(b) Reconstructed Shower $\cos(\theta)$ 

(c) Maximum PSV score between Plane 0 and Plane 2

FIG. 14: Distributions of the sideband sample in: (a) reconstructed shower energy; (b) reconstructed shower $\cos(\theta)$; and (c) the maximum PSV score between Plane 0 and Plane 2. A mild data deficit is observed in the highest PSV score bin, but data is found to be consistent with prediction within systematic uncertainty over all three variables.

B. Coherent-Rich Selection Results

For the single-photon selection, 70 events are observed in data compared to an MC prediction of 55.8, leading to a data-MC ratio of $1.25 \pm 0.32(\text{sys})$. Distributions of the reconstructed shower energy, angle ($\cos\theta$) and the maximum PSV score on planes 0 and 2 were previously shown in Fig. 9. Overall, the agreement between data-MC is good according to the goodness-of-fit tests with p -values of 0.251, 0.830, and 0.476, respectively.






FIG. 15: Comparison of observed data and MC CV predictions for the coherent-rich selection. The red (green) line and surrounding band represent the MC prediction and associated systematic uncertainty before (after) the 2γ constraint is applied. The breakdown of MC prediction in different categories can be found in Tab. II.

For the coherent-rich subset selection, 34 data events are observed, compared to a constrained prediction of $29.0 \pm 5.2(\text{syst}) \pm 5.4(\text{stats})$ and an unconstrained prediction of $34.0 \pm 8.4(\text{syst}) \pm 5.8(\text{stats})$. The constrained background prediction is lower than unconstrained due to the mild deficit observed in data in the NC π^0 high statistics two-shower constraining channels, as can be seen in Fig. 13. The comparison between data and MC together with results from goodness-of-fit tests are shown in Fig 15. Data-MC agreement slightly worsens after the constraint from the 2γ channels is applied, from a p -value of 1.00 to a p -value of 0.51, but nonetheless remains well within expectations. Distributions of the 34 data events and our simulation prediction can be found in Fig. 16.

The observed data in the coherent-rich selection is fit to extract a normalization scaling factor x for the NC coherent 1γ process, with full systematic uncertainties included except for the GENIE cross section uncertainty for the signal process. The best-fit for x is found to be $x = 6.20$ with a $\chi^2/\text{ndf} = 9.7/15$. Figure 17 shows the confidence level (C.L.) for different values of x as evaluated using the Feldman-Cousins (FC) frequentist method [64], with the C.L. under the assumption of Wilks' theorem overlaid. We see that curves from FC and Wilks' theorem agree with each other in the ma-



(a) Reconstructed Shower Energy



(b) Reconstructed Shower $\cos(\theta)$

FIG. 16: Distributions of (a) reconstructed shower energy; (b) reconstructed shower $\cos(\theta)$ for all events in the coherent-rich subset selection. We note that, due to low statistics, the fit only considers the total number of events in these distributions; the distributions are shown as additional information.

majority of the parameter space and start to differ in the parameter region near the boundary, as expected. Given the observed data, the bound at 90% C.L. is placed at $x = 24.0$ and the 95% C.L. bound is at $x = 28.0$ for the NC coherent 1γ , corresponding to a limit on the flux-averaged cross section of $< 1.49 \times 10^{-41} \text{ cm}^2$ at 90% C.L.. Note that the slight fluctuations in the confidence level curve for the FC method are due to the finite number of pseudo-experiments used.




FIG. 17: The confidence level for the NC coherent 1γ normalization scaling factor x derived from observed data in the coherent-rich selection. The C.L. evaluated with the Feldman-Cousins method is shown in pink with C.L. assuming the validity of Wilks’ theorem overlaid in blue. Dotted, dashed and large-dashed horizontal lines correspond to 2σ , 90% and 1σ respectively.

VIII. SUMMARY

In this paper, we present the world’s first search for the neutrino-induced NC coherent single-photon process with a single shower topology and no other visible vertex activity. The search is performed with the first three years of data from the MicroBooNE detector, with a neutrino beam of $\langle E_\nu \rangle \sim 0.8$ GeV from the Fermilab Booster Neutrino Beam. Differences in the shower kinematic properties between the shower originating from coherent photons and that from other backgrounds are leveraged to yield effective background removal, leading to 99.97% and $> 95\%$ rejection efficiency on the cosmic and NC non-coherent $1\pi^0$ background relative to Pandora’s topological selection, respectively. New tools utilizing low-level TPC charge information are developed to identify and reject low-energy proton tracks near the vertex, which further reduces the dominant NC non-coherent $1\pi^0$ background by 48%. This search yields a signal-to-background ratio of $\approx 1 : 30$ in the final coherent-rich selection, and good agreement between data and prediction

is observed. This leads to the world’s first experimental upper limit on the cross section of neutrino-induced NC coherent single-photon production on an argon nucleus below 1 GeV, of $\sigma < 1.49 \times 10^{-41} \text{cm}^2$ at 90% C.L., corresponding to 24.0 times the prediction in Ref. [13].

The proton veto tool developed in this analysis has shown great potential in rejecting low-energy protons and represents the first high-level use of MeV scale energy deposits in MicroBooNE. This veto can be easily adapted to other coherent interaction searches. Furthermore, it indicates the way forward for enhancing the reconstruction of low-energy particles, including searching for evidence of activity due to neutron reinteractions. While this search has limited sensitivity, the selections developed can be applied to future LArTPC experiments for a search for this process with higher statistics. An example of this is the upcoming SBND experiment which is expected to collect $\mathcal{O}(20)$ times more data [20] than MicroBooNE, potentially enabling a limit on this process much closer to the SM-predicted rate.

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, Office of High Energy Physics HEP User Facility. Fermilab is managed by Fermi Forward Discovery Group, LLC, acting under Contract No. 89243024CSC000002. 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); the UK Research and Innovation (UKRI) Future Leaders Fellowship; and the NSF AI Institute for Artificial Intelligence and Fundamental Interactions. Additional support for the laser calibration system and cosmic ray tagger was provided by the Albert Einstein Center for Fundamental Physics, Bern, Switzerland. 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. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) public copyright license to any Author Accepted Manuscript version arising from this submission.

[1] R. L. Workman *et al.* (Particle Data Group), PTEP **2022**, 083C01 (2022).

[2] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **106**, L051102 (2022).

- [3] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **107**, 012004 (2023).
- [4] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **104**, 052002 (2021).
- [5] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **105**, L051102 (2022).
- [6] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. Lett. **125**, 201803 (2020).
- [7] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **102**, 112013 (2020).
- [8] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. Lett. **123**, 131801 (2019).
- [9] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Phys. Rev. D **103**, 052002 (2021).
- [10] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. Lett. **130**, 011801 (2023), arXiv:2210.10216 [hep-ex].
- [11] P. A. Machado, O. Palamara, and D. W. Schmitz, Annu. Rev. Nucl. Part. Sci. **69**, 363 (2019), <https://doi.org/10.1146/annurev-nucl-101917-020949>.
- [12] A. Abud Abed *et al.* (DUNE), Phys. Rev. D **105**, 072006 (2022), arXiv:2109.01304 [hep-ex].
- [13] E. Wang, L. Alvarez-Ruso, and J. Nieves, Phys. Rev. C **89**, 015503 (2014).
- [14] E. Wang, L. Alvarez-Ruso, and J. Nieves, Phys. Lett. B **740**, 16 (2015), arXiv:1407.6060 [hep-ph].
- [15] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. Lett. **128**, 111801 (2022).
- [16] S. N. Gninenko, Phys. Rev. Lett. **103**, 241802 (2009), arXiv:0902.3802 [hep-ph].
- [17] O. Fischer, A. Hernández-Cabezudo, and T. Schwetz, Phys. Rev. D **101**, 075045 (2020), arXiv:1909.09561 [hep-ph].
- [18] E. Bertuzzo, S. Jana, P. A. N. Machado, and R. Zukanovich Funchal, Phys. Rev. Lett. **121**, 241801 (2018), arXiv:1807.09877 [hep-ph].
- [19] P. Ballett, S. Pascoli, and M. Ross-Lonergan, Phys. Rev. D **99**, 071701 (2019), arXiv:1808.02915 [hep-ph].
- [20] C. Adams *et al.* (LArTPC), in *Snowmass 2013: Snowmass on the Mississippi* (2013) arXiv:1309.7987 [physics.ins-det].
- [21] P. Abratenko *et al.* (MicroBooNE Collaboration), “Enhanced Search for Neutral Current Δ Radiative Single-Photon Production in MicroBooNE,” (2025), <https://microboone.fnal.gov/nc-delta-2025/>.
- [22] P. Abratenko *et al.* (MicroBooNE Collaboration), “Inclusive Search for Anomalous Single-Photon Production in MicroBooNE,” (2025), <https://microboone.fnal.gov/inclusive-single-photon-2025/>.
- [23] D. R. Nygren, 1974 PEP summer study, eConf **C740805**, 58 (1974).
- [24] C. Rubbia, CERN-EP/77-08 (1977).
- [25] R. Acciarri *et al.* (ArgoNeuT Collaboration), Phys. Rev. D **99**, 012002 (2019).
- [26] S. Andringa *et al.*, J. Phys. G **50**, 033001 (2023).
- [27] W. Castiglioni, W. Foreman, I. Lepetic, B. R. Littlejohn, M. Malaker, and A. Mastbaum, Phys. Rev. D **102**, 092010 (2020), arXiv:2006.14675 [physics.ins-det].
- [28] P. Abratenko *et al.* (MicroBooNE Collaboration), JINST **17**, P11022 (2022), arXiv:2203.10147 [physics.ins-det].
- [29] P. Abratenko *et al.* (MicroBooNE), Phys. Rev. D **109**, 052007 (2024), arXiv:2307.03102 [hep-ex].
- [30] R. Acciarri *et al.* (MicroBooNE Collaboration), JINST **12**, P02017 (2017), arXiv:1612.05824 [physics.ins-det].
- [31] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Phys. Rev. D **79**, 072002 (2009).
- [32] C. Andreopoulos *et al.*, Nucl. Instrum. Meth. A **614**, 87 (2010).
- [33] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **105**, 072001 (2022).
- [34] K. Abe *et al.* (T2K Collaboration), Phys. Rev. D **93**, 112012 (2016).
- [35] C. Berger and L. M. Sehgal, Phys. Rev. D **79**, 053003 (2009).
- [36] A. Bodek, S. Avvakumov, R. Bradford, and H. S. Budd, Eur. Phys. J. C **53**, 349 (2008), arXiv:0708.1946 [hep-ex].
- [37] R. Merenyi, W. A. Mann, T. Kafka, W. Leeson, B. Saitta, J. Schneps, M. Derrick, and B. Musgrave, Phys. Rev. D **45**, 743 (1992).
- [38] E. Snider and G. Petrillo, Journal of Physics: Conference Series **898**, 042057 (2017).
- [39] S. Agostinelli *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **506**, 250 (2003).
- [40] C. Adams *et al.* (MicroBooNE Collaboration), JINST **13**, P07006 (2018), arXiv:1802.08709 [physics.ins-det].
- [41] C. Adams *et al.* (MicroBooNE Collaboration), JINST **13**, P07007 (2018), arXiv:1804.02583 [physics.ins-det].
- [42] M. Jaskolski and M. Wojcik, J. Phys. Chem. A **115**, 4317 (2011), PMID: 21473614, <https://doi.org/10.1021/jp201149w>.
- [43] C. Adams *et al.* (MicroBooNE Collaboration), JINST **15**, P03022 (2020), arXiv:1907.11736 [physics.ins-det].
- [44] C. Adams *et al.* (MicroBooNE Collaboration), JINST **15**, P07010 (2020), arXiv:1910.01430 [physics.ins-det].
- [45] C. Adams *et al.* (MicroBooNE Collaboration), JINST **16**, P04004 (2021), arXiv:2012.14324 [physics.ins-det].
- [46] R. Acciarri *et al.* (MicroBooNE Collaboration), Eur. Phys. J. C **78**, 82 (2018), arXiv:1708.03135 [hep-ex].
- [47] R. Frühwirth, Nucl. Instrum. Meth. A **262**, 444 (1987).
- [48] R. Acciarri *et al.* (MicroBooNE), JINST **12**, P09014 (2017), arXiv:1704.02927 [physics.ins-det].
- [49] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **107**, 012004 (2023).
- [50] L. Alvarez-Ruso *et al.*, Eur. Phys. J. ST **230**, 4449 (2021).
- [51] P. Abratenko *et al.* (MicroBooNE), Phys. Rev. Applied **15**, 064071 (2021), arXiv:2101.05076 [physics.ins-det].
- [52] T. Chen and C. Guestrin, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16 (ACM, 2016).
- [53] M. Ester, H. P. Kriegel, J. Sander, and X. Xiaowei, Data Min. Knowl. Discov , 169–194 (1996).
- [54] D. H. Wright and M. H. Kelsey, Nucl. Instrum. Meth. A **804**, 175 (2015).
- [55] J. Calcutt, C. Thorpe, K. Mahn, and L. Fields, JINST **16**, P08042 (2021).
- [56] P. Abratenko *et al.*, Eur. Phys. J. C **82**, 454 (2022).
- [57] P. Abratenko *et al.* (MicroBooNE Collaboration), JINST **15**, P12037 (2020), arXiv:2008.09765 [physics.ins-det].
- [58] R. Acciarri *et al.* (ArgoNeuT), JINST **8**, P08005 (2013), arXiv:1306.1712 [physics.ins-det].
- [59] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **105**, 112005 (2022).
- [60] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **105**, 112004 (2022).
- [61] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys. Rev. D **105**, 112003 (2022).
- [62] P. Abratenko *et al.* (MicroBooNE Collaboration), Phys.

Rev. Lett. **128**, 241801 (2022).
[63] X. Ji, W. Gu, X. Qian, H. Wei, and C. Zhang, Nucl. Instrum. Meth. A **961**, 163677 (2020).
[64] G. J. Feldman and R. D. Cousins, Phys. Rev. D **57**, 3873 (1998), arXiv:physics/9711021.