

Experimental and Phenomenological Investigations of the MiniBooNE Anomaly

by

Nicholas Kamp

B.S.E., University of Michigan (2019)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

© 2023 Nicholas Kamp. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by

Department of Physics
May 23, 2023

Certified by

Janet M. Conrad
Professor
Thesis Supervisor

Accepted by

Lindley Winslow
Associate Department Head of Physics

Experimental and Phenomenological Investigations of the MiniBooNE Anomaly

by

Nicholas Kamp

Submitted to the Department of Physics
on May 23, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

The 4.8σ excess of electron neutrino-like events reported by the MiniBooNE experiment at Fermilab’s Booster Neutrino Beam (BNB) is one of the most significant and longest standing anomalies in particle physics. This thesis covers a range of experimental and theoretical efforts to elucidate the origin of the MiniBooNE low energy excess (LEE). We begin with the follow-up MicroBooNE experiment, which took data along the BNB from 2016 to 2021. The detailed images produced by the MicroBooNE liquid argon time projection chamber enable a suite of measurements that each test a different potential source of the MiniBooNE anomaly. This thesis specifically presents MicroBooNE’s search for ν_e charged-current quasi-elastic (CCQE) interactions consistent with two-body scattering. The two-body CCQE analysis uses a novel reconstruction process, including a number of deep-learning based algorithms, to isolate a sample of ν_e CCQE interaction candidates with 75% purity. The analysis rules out an entirely ν_e -based explanation of the MiniBooNE excess at the 2.4σ confidence level. We next perform a combined fit of MicroBooNE and MiniBooNE data to the popular $3 + 1$ model; even after the MicroBooNE results, allowed regions in Δm^2 - $\sin^2 2\theta_{\mu e}$ parameter space exist at the 3σ confidence level. This thesis also demonstrates that, due to nuclear effects in the low-energy cross section behavior, the MicroBooNE data are consistent with a $\bar{\nu}_e$ -based explanation of the MiniBooNE LEE at the $< 2\sigma$ confidence level. Next, we investigate a phenomenological explanation of the MiniBooNE excess involving both an eV-scale sterile neutrino and a dipole-coupled MeV-scale heavy neutral lepton (HNL). It is shown that a 500 MeV HNL can accommodate the energy and angular distributions of the LEE at the 2σ confidence level while avoiding stringent constraints derived from MINER ν A elastic scattering data. Finally, we discuss the Coherent CAPTAIN-Mills (CCM) experiment—a 10-ton light-based liquid argon detector at Los Alamos National Laboratory. The background rejection achieved from a novel Cherenkov-based reconstruction algorithm will give CCM world-leading sensitivity to a number of beyond-the-Standard Model physics scenarios, including dipole-coupled HNLs.

Thesis Supervisor: Janet M. Conrad
Title: Professor

Acknowledgments

Development from a starry-eyed first year graduate student into a competent researcher is like the MSW effect: it doesn't happen in a vacuum. There are many people who have helped me along the way in both physics and life, without whom I would never have gotten to the point of writing this thesis.

First and foremost, I owe an immense debt of gratitude to Janet Conrad. Janet has been an incredible mentor to me during my time as a graduate student; her advice and wisdom have helped me become the scientist I wanted to be when I came to MIT four years ago. I'd like to specifically thank Janet for taking my interests seriously and working with me to develop research projects that matched them. Creativity, ingenuity, enthusiasm, and kindness run rampant in the Conrad group—I will always be grateful for being offered a spot in it. I look forward to many years of fruitful collaboration to come.

To my partner, Wenzer: thank you for the love, support, and patience over the last two years. Life is not so hard when we can act as a restoring force for one another—I am grateful for having been able to rely on it while writing this thesis. I look forward with great excitement to our future adventures together.

Thank you to my past mentors: Christine Aidala for introducing me to particle physics through POLARIS, Robert Cooper for asking me about my research at APS DNP 2016, Bill Louis for answering my many questions about neutrino physics in my first summer at LANL, Richard Van de Water for teaching me to follow the data (which greatly influenced my choice of graduate research), and Josh Spitz for helping develop confidence as a neutrino physicist on JSNS². I'd also like to thank Christopher Mauger and the rest of the Mini-CAPTAIN team for an introduction to what it takes to run a particle physics experiment at an accelerator.

Thank you to members of the Conrad group past and present: those I worked with, Lauren Yates, Adrian Hourlier, Jarrett Moon, Austin Schneider, Darcy Newmark, Alejandro Diaz, and John Hardin, for being fantastic collaborators, and those I did not, Loyd Waits, Joe Smolsky, Daniel Winklehner, Philip Weigel, and Josh

Villareal, for making MIT a brighter place. Thank you especially to Austin for your infinite patience in answering my many questions on statistics and software—each of our projects has been a great pleasure.

To my MicroBooNE collaborators not mentioned above, Taritree Wongjirad, Katie Mason, Joshua Mills, Polina Abratenko, Ran Itay, Mike Shaevitz, Georgia Karagiorgi, Davio Cianci, Rui An, and everyone else: thank you for your excellent teamwork in putting together the Deep Learning analysis.

To my CCM collaborators not mentioned above, Edward Dunton, Mayank Tripathi, Adrian Thompson, Will Thopmson, Marisol Chávez Estrada, and everyone else: thank you for the invigorating research and discussion over the last couple of years, and best of luck with CCM200!

Thank you to Carlos Argüelles, Mike Shaevitz, Matheus Hostert, Stefano Vergani, and Melissa Uchida for your excellent mentorship and collaboration in our phenomenological endeavors together. Thank you specifically to Carlos for giving me a welcome introduction to neutrino phenomenology, Mike for ensuring the robustness of each analysis, and Matheus for patiently answering my many model-building questions.

To all of my friends not mentioned above; Jack, Ryan, Alexis, Melissa, Ben, Bhaamati, Charlie, Vincent, Patrick, Caolan, Ouail, Artur, Sam, Zhiquan, Rebecca, Felix, Lila, Rahul, Brandon, Field, Kelsey, Woody, Joey, Rory, Cooper, Daniel, Kaliröë, Elena, and everyone else: thank you for all of the great memories—climbing, hiking, skiing, playing music, eating, drinking, commiserating, and laughing—over the past four years. Thank you especially to the last three for making preparation for the oral exam significantly more enjoyable.

Thank you to the MIT administrative staff, including (but not limited to) Lauren Saragosa, Karen Dow, Catherine Modica, Sydney Miller, Alisa Cabral, and Elsy Luc, for helping make graduate school a more manageable endeavor. Thank you also to the rest of my thesis committee, Joseph Formaggio and Washington Taylor, for helping me get through the final part of graduate school.

Finally, thank you to my parents, Jim and Carla, my siblings, Serafina and Daniel,

and the rest of my family. From elementary school science fairs to Saturday Morning Physics to today, none of this would have been possible without your love and support.

Contents

1	Introduction	29
1.1	A Brief History of the Neutrino	29
1.2	Neutrinos in the Standard Model	35
1.3	Massive Neutrinos	39
1.4	Anomalies in the Neutrino Sector	48
2	The MiniBooNE Experiment	59
2.1	Overview of MiniBooNE	59
2.1.1	The Booster Neutrino Beam	60
2.1.2	The MiniBooNE Detector	61
2.2	The MiniBooNE Low Energy Electron-Like Excess	63
3	The MicroBooNE Detector	71
3.1	Liquid Argon Time Projection Chamber	71
3.1.1	Cryogenics	73
3.1.2	LArTPC Drift System	74
3.1.3	Light Collection System	75
3.2	TPC Signal Processing	79
3.2.1	Noise Filtering	80
3.2.2	Deconvolution	81
4	The MicroBooNE Electron Neutrino Analysis: Overview and Selection	85

4.1	Dataset and Simulation	86
4.2	The Electron Low Energy Excess Template	88
4.3	Philosophy Behind the Two-Body CCQE Analysis	90
4.4	Reconstruction	93
4.4.1	Convolutional Neural Networks in LArTPCs	96
4.4.2	Vertex and Track Reconstruction	98
4.4.3	Publication: <i>Electromagnetic shower reconstruction and energy validation with Michel electrons and π^0 samples for the deep-learning-based analyses in MicroBooNE</i>	102
4.5	1e1p Event Selection	132
4.5.1	Basic Data Selection Criteria	132
4.5.2	Boosted Decision Tree Ensemble	133
4.5.3	Particle Identification Cuts	137
4.5.4	The Final 1e1p Sample	139
5	The MicroBooNE Electron Neutrino Analysis: Results and Discussion	145
5.1	First Results from the Two-Body CCQE Analysis	146
5.1.1	Background Estimation	147
5.1.2	Evaluation of Systematic Uncertainties	151
5.1.3	Constraint from the $1\mu 1p$ Sample	156
5.1.4	Blinded Analysis Approach	158
5.2	Statistical Interpretation	159
5.2.1	Goodness of Fit	161
5.2.2	Two Hypothesis Test	162
5.2.3	Signal Strength Scaling Test	164
5.3	Discussion and Outlook	164
5.3.1	Publication: <i>MiniBooNE and MicroBooNE Combined Fit to a $3 + 1$ Sterile Neutrino Scenario</i>	167

5.3.2	Publication: <i>Implications of MicroBooNEs low sensitivity to electron antineutrino interactions in the search for the Mini-BooNE excess</i>	178
6	Neutrissimos: Heavy Neutral Leptons with a Dipole Moment	195
6.1	Dipole-Portal Neutrissimos	196
6.2	Overview of the Mixed Model	198
6.3	Neutrissimos in MiniBooNE	201
6.3.1	Simulation in LeptonInjector	204
6.3.2	Fits to the MiniBooNE Excess	207
6.4	Publication: <i>Dipole-coupled neutrissimo explanations of the MiniBooNE excess including constraints from MINERvA data</i>	211
7	The Coherent CAPTAIN-Mills Experiment	229
7.1	The CCM Beamline and Detector	230
7.2	Cherenkov Light Reconstruction	237
7.2.1	Simulation-Based Sensitivity Estimation	240
7.2.2	Identifying Cherenkov Light in Data	243
7.3	Neutrissimos in CCM	248
8	Conclusions and Future Prospects	259
A	Publication: <i>Convolutional neural networks for shower energy prediction in liquid argon time projection chambers</i>	265
B	Signal Event Displays from the Two-Body CCQE Analysis	293

List of Figures

1-1	Telegram from Fred Reines and Clyde Cowan informing Wolfgang Pauli of their detection of neutrinos from a nuclear reactor.	30
1-2	The deficit of the observed solar ν_e flux compared with the theoretical expectation. The Homestake experiment is shown on the far left; follow-up solar neutrino measurements confirming the deficit are also shown, including the 2002 SNO result which brought forth a solution to the solar neutrino problem. Figure from Ref. [1].	31
1-3	The up-down asymmetry measured in SuperK as a function of lepton momentum, separated into e -like and μ -like events as well as fully-contained (FC) and partially-contained (PC) events. The dashed line indicates the best fit to $\nu_\mu \rightarrow \nu_\tau$ oscillations. Figure from Ref. [2]. . .	34
1-4	Measurement of the solar ^8B flux from the SNO collaboration, broken down into the ν_e and $\nu_{\mu,\tau}$ sub-components. Measurement of the CC, NC, and ES channels show up as slices in the two-dimensional flux parameter space. Figure from Ref. [3].	34
1-5	Diagrams contributing to νe^- elastic scattering	39
1-6	Diagrams contributing to $\bar{\nu} e^-$ elastic scattering	39
1-7	Diagrams contributing to neutrino-nucleon charged-current quasielastic scattering	40
1-8	CC inclusive neutrino and antineutrino nucleon scattering cross sections as a function of neutrino energy. Figure from Ref. [4].	40

1-9	The LSND excess of $\bar{\nu}_e$ events on top of the predicted SM background (green and red regions). The blue region indicates the best fit to $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations via a sterile neutrino state. Figure from Ref. [5]	49
1-10	The MiniBooNE electron-like channel data and SM background prediction for the entire neutrino mode dataset, as a function of the reconstructed neutrino energy.	51
1-11	Data contributing to the reactor antineutrino anomaly, indicating the $\sim 5\%$ flux deficit observed by short-baseline reactor neutrino experiments. The red line indicates the prediction incorporating SM neutrino oscillations only, while the blue line shows an example prediction including a sterile neutrino. Figure from Ref. [6].	52
1-12	Data contributing to the gallium anomaly, indicating the $\sim 20\%$ deficit in the ^{71}Ge production rate observed by SAGE, GALLEX, and BEST. Figure from Ref. [7].	53
1-13	Preferred regions in $\sin^2 2\theta_{ee}-\Delta m^2$ parameter space to explain the RAA [8] (green contour) and gallium anomaly [9] (blue regions). The total excluded region from other experiments (grey region) is also shown. Figure from Ref. [9].	56
1-14	Preferred regions in $\sin^2 2\theta_{\mu e}-\Delta m^2$ parameter space to explain the LSND anomaly [5] (filled contours) and MiniBooNE anomaly [10] (open contours). Figure from Ref. [10].	57
1-15	Graphical representation of the tension observed in 3+1 global fits between different subsets of the experimental landscape. Figure 1-15a shows the tension between ν_e appearance experiments and ν_e/ν_μ disappearance experiments observed in Ref. [11]. Figure 1-15b shows the tension between allowed regions from ν_e appearance (lower right), ν_e disappearance (upper right), and ν_μ disappearance (upper left) experiments observed in Ref. [12], which includes the latest results from the BEST experiment.	58

2-1	A schematic depiction of the BNB at Fermilab, including the downstream MiniBooNE detector. Figure from Ref. [13].	61
2-2	Breakdown of the neutrino flux at the BNB in neutrino (left) and antineutrino (right) mode. Figure from Ref. [14]	61
2-3	The MiniBooNE detector situated in the cylindrical detector hall (left) and an image of the interior of the MiniBooNE detector (right), showing the PMTs in both the signal and veto regions. Figure from Ref. [15].	64
2-4	Visual representations of particle identification in MiniBooNE. Figure 2-4a shows a schematic representation of the detector signature from the three main particle classes in MiniBooNE: muons, electrons, and neutral pions. Figure 2-4b shows the MiniBooNE log-likelihood-ratio between the e -like and μ -like hypothesis as a function of reconstructed neutrino energy, considering both simulated ν_e CCQE (top) and ν_μ CCQE (bottom) interactions.	64
2-5	The E_ν^{QE} distribution of the MiniBooNE e -like excess in the total neutrino mode (figure 2-5a) and antineutrino mode (figure 2-5b) datasets. The observation and SM prediction in each bin are shown by the data points and colored histograms, respectively.	68
2-6	The lepton visible energy (figure 2-6a) and $\cos\theta$ (figure 2-6b) distributions of the MiniBooNE e -like excess in the total neutrino mode dataset. The observation and SM prediction in each bin are shown by the data points and colored histograms, respectively. Figures from Ref. [10].	68
2-7	The MiniBooNE and LSND excesses as a function of the ratio L/E . The MiniBooNE data is separated into neutrino and antineutrino mode. Figure from Ref. [10].	70

3-1	Schematic depictions of the MicroBooNE LArTPC. Figure 3-1a shows the detection process for charged particles from a neutrino interaction in a MicroBooNE-like LArTPC. Figure 3-1b shows a cross-sectional view of the MicroBooNE detector along the $-\hat{z}$ direction. Figures from Ref. [16].	73
3-2	Figure 3-2a shows a close-up image of the cathode plane of the MicroBooNE LArTPC. The stainless steel field cage tubes can also be seen surrounding the active volume. Figure 3-2b shows a cross-sectional map of the electric field at the edge of the active volume, considering a cathode plane voltage of -128 kV. The legend shows the field strength in units of V/m. Figures from Ref. [16].	76
3-3	Figure 3-3a shows a photograph of a single wire carrier board with 32 mounted wires. Figure 3-3b shows the fully-assembled MicroBooNE LArTPC, highlighting the anode plane mounted on the stainless steel frame. Figures from Ref. [16].	76
3-4	The LAr scintillation light spectrum, TPB ultra-violet absorption spectrum, TPB emission spectrum, PMT quantum efficiency, and PMT surface transmission efficiency as a function of photon wavelength. Figure from Ref. [17].	79
3-5	Figure 3-5a shows a photograph of a single PMT system used in the MicroBooNE detector. The acrylic window here has not yet been coated with TPB. Figure 3-5b shows the PMT signal from a stopped cosmic muon (top) that decays to a Michel (bottom). Figures from Ref. [16].	79
3-6	2D displays of the signal from wires in one of the induction planes in a single data event, before and after the application of the offline noise filters. Figure from Ref. [18].	81
3-7	A U plane event display of a candidate neutrino interaction in the MicroBooNE data. The impact of the 1D and 2D deconvolution algorithms on the post-noise-filtering signal is shown. Figure from Ref. [19].	84

4-1	Figure 4-1a shows the unfolded ν_e prediction in MiniBooNE calculated using the first 6.46×10^{20} POT of neutrino mode data [20]. Figure 4-1b shows the unfolded eLEE model weights derived from the first 12.84×10^{20} POT of MiniBooNE neutrino mode data, which constitute the MicroBooNE eLEE model.	90
4-2	The expected distribution of ν_e interactions in MicroBooNE as a function of the true ν_e energy. The dotted line shows the expectation from the MicroBooNE eLEE model of the MiniBooNE excess discussed in section 4.2.	94
4-3	An example candidate $1e1p$ data event in MicroBooNE, including the raw LArTPC collection plane image (left) and the pixel labels assigned from SparseSSNet (right).	94
4-4	The relationship between the final state lepton (left) and proton (right) energy and scattering angle, for different neutrino energies, in ν_μ (top) and ν_e (bottom) CCQE scattering.	95
4-5	Figure 4-5a shows a diagram of the SparseSSNet U-ResNet architecture. Figure 4-5b shows the SparseSSNet pixel labels on a simulated image in the training dataset. Figures from Ref. [21].	99
4-6	Figure 4-6a shows a diagram of the MPID architecture. Figure 4-6b shows the MPID image scores on an example simulated event. Figures from Ref. [22].	100
4-7	The fraction of EM showers reconstructed to within 5% of the true deposited energy as a function of the fraction of unresponsive wires in the LArTPC collection plane, considering three different methods. The ResNet and Inception network significantly outperform the traditional linear calibration between charge and energy. Figure from Ref. [23].	101

4-8	Figure 4-8a shows the angular metric that is minimized to find a 3D neutrino vertex candidate. Figure 4-8b shows the iterative track reconstruction algorithm, which relies on calculating distances (L_1 and L_2) and angles (θ and ϕ) with respect to the previous point and the end of the track. Figures from Ref. [24].	101
4-9	An example of an event that fails the shower energy consistency cut, because the EM shower passes through an unresponsive region of the collection plane.	133
4-10	Figure 4-10a shows the distribution of the fractional shower energy consistency variable for events with an old $1e1p$ BDT score above/below 0.7. Figure 4-10b shows the efficiency with which events above/below the old $1e1p$ BDT cutoff pass the fractional consistency cut as a function of the chosen upper bound.	134
4-11	The F score of each variable for one of the $1e1p$ BDTs in the ensemble from run 1, run2, and run 3.	137
4-12	Figure 4-12a and figure 4-12b show the MC distribution of the $1e1p$ BDT ensemble average score over all three run periods for the full [0,1] range and zoomed in to the [0.95,1] range, respectively.	137
4-13	Figure 4-13a and figure 4-13b show the predicted ν_e event rate in run period 2 and run period 3, respectively, using both the run period 2 ensemble and the run period 3 ensemble.	138
4-14	Figure 4-14a and figure 4-14b show the fractional difference in average BDT score $(S_n - S_0)/S_0$ as a function of the number of omitted BDTs n over the simulation from run period 2 and run period 3, respectively. The red histogram shows the actual distribution of the number of omitted BDTs over the run period 2 and run period 3 simulation samples, respectively. Scores are calculated using the run period 3 and run period 2 BDT ensemble, respectively.	138

4-15	Figure 4-15a and figure 4-15b show the MPID electron and muon score, respectively, as a function of the reconstructed electron energy in intrinsic ν_e MC events.	139
4-16	The E_ν^{range} distribution for the $1e1p$ signal sample, showing only the predicted event rate from the MC. The prediction from the eLEE model is shown in the dashed blue line.	140
4-17	Figure 4-17a shows the distribution of the fractional error on the neutrino energy for MC events in the $1e1p$ signal sample, restricted to $200 < E_\nu^{\text{Range}}$ [MeV] < 1200 . Figure 4-17b shows the 2D distribution of fractional error as a function of the true neutrino energy.	140
4-18	Figure 4-18a shows the post-vertex-identification efficiency of true ν_e CCQE selection for subsequent stages of the $1e1p$ cuts. Figure 4-18b shows the true ν_e CCQE event rates over the full run 1-3 dataset after subsequent stages of the $1e1p$ cuts.	141
5-1	Top: pixel intensity (color scale is in PIU as defined in section 4.4); Bottom: SparseSSNet labels; Left to Right: U, V, Y, planes. The white circle indicates the reconstructed vertex. The horizontal axis corresponds to the wire plane direction and the vertical axis corresponds to the electron drift direction, which is measured using the arrival time of charge on the wires.	147
5-2	The $1e1p$ sample E_ν distribution, comparing data (black points) to the unconstrained prediction (stacked histogram) in the $200 < E_\nu < 1200$ MeV region. The eLEE model prediction is represented by the dashed blue line. The prediction is presented in terms of both interaction type (figure 5-2a) and final state topology (figure 5-2b).	148
5-3	Average $1e1p$ BDT ensemble score distribution comparing data to the unconstrained prediction.	148

5-4	Comparison between data and unconstrained prediction in the E_e (figure 5-4a), E_p (figure 5-4b), θ_e (figure 5-4c), and $E_\nu^{QE-\ell}$ (figure 5-4d) distributions of the $1e1p$ sample.	149
5-5	The fit to the ν_μ background distribution to the $1e1p$ analysis. The shape fit is performed at a loose BDT score cutoff of 0.7 (figure 5-5a) and scaled to the signal cutoff of 0.95 (figure 5-5b). Blue points represent the prediction from the simulation, with error bars representing the Gaussian approximation of the statistical error (quadrature sum of event weights). The orange line and corresponding shaded region represent prediction and uncertainty, respectively, coming from the Landau+linear fit.	152
5-6	The data and MC prediction for events with a $1e1p$ BDT score inside $[0.7, 0.95]$. Good agreement is observed between data and prediction. The prediction incorporating the Landau+linear background fit is shown by the red line.	153
5-7	The uncertainty in each bin of the E_ν^{range} distribution of the $1e1p$ (figure 5-7a) and $1\mu1p$ (figure 5-7b) samples.	155
5-8	The joint covariance (figure 5-8a) and correlation (figure 5-8b) matrices for the E_ν^{range} distribution of the $1e1p$ and $1\mu1p$ samples.	156
5-9	The E_ν^{range} distribution in the $1\mu1p$ channel, comparing data to the MC prediction.	157
5-10	Fractional systematic uncertainty in the $1e1p$ E_ν^{range} distribution before and after the $1\mu1p$ constraint.	157
5-11	Comparison between data and prediction in the $1e1p$ E_ν^{range} distribution after applying the $1\mu1p$ constraint procedure.	158
5-12	Comparison between data and prediction in the $1e1p$ BDT ensemble average score distribution within the range $[0.7, 0.95]$	159
5-13	Distributions of the $\Delta\chi^2$ test statistic defined in equation (5.7) for H_0 (red) and H_1 (blue), calculated by generating 10^5 pseudo-experiments under each hypothesis. The $\Delta\chi^2$ value of the data is also shown.	163

5-14	Confidence intervals on x_{LEE} calculating using the Feldman-Cousins procedure. The solid and dotted lines indicate the confidence level with which a given x_{LEE} is disfavored, calculated using the Feldman-Cousins method [25] and Wilks theorem [26], respectively. The MiniBooNE statistical and systematic errors are shown as a band around $x_{\text{LEE}} = 1$.	165
5-15	Figure 5-15a shows the observation compared to the nominal (H_0) prediction in all four signal channels from the three MicroBooNE ν_e analyses, including statistical errors on the data points and systematic errors on the prediction. The eLEE prediction (H_1) is also indicated by the red line. Figure 5-15b shows the observed 1σ and 2σ confidence intervals on x_{LEE} from all four signal channels. The 2σ expected sensitivity of each channel is shown in red.	166
6-1	Feynman diagram depicting the effective dipole operator of equation (6.1).	198
6-2	New interactions involving the neutrissimo that are enabled by the dipole operator in equation (6.1), including three-body π^0 decay (figure 6-2a), Primakoff-like upscattering (figure 6-2b), and neutrissimo decay (figure 6-2c).	199
6-3	3+1 global fits including MiniBooNE, considering global (left), appearance-only (middle), and disappearance-only (right) experiments. The allowed regions in 3 + 1 parameter space at the 90%, 95%, and 99% confidence levels are shown by the red, green, and blue points, respectively. The best-fit point is indicated by the star.	201
6-4	3+1 global fits without MiniBooNE, considering global (left), appearance-only (middle), and disappearance-only (right) experiments. The allowed regions in 3 + 1 parameter space at the 90%, 95%, and 99% confidence levels are shown by the red, green, and blue points, respectively. The best-fit point is indicated by the star.	201

6-5	Schematic depiction of the neutrissimo model in MiniBooNE as simulated using <code>LeptonInjector</code> . Figure 6-5a shows the simulation of Primkaoff upscattering along the beamline, and figure 6-5b shows an example of upscattering, neutrissimo decay and pair-production within the MiniBooNE detector.	207
6-6	Allowed regions at the 95% and 3σ confidence level in $d_{\mu\mathcal{N}}-m_{\mathcal{N}}$ obtained through fits to the MiniBooNE excess in the E_{ν}^{QE} and $\cos\theta$ distributions. Existing 2σ constraints on this model are indicated by the grey regions.	210
6-7	Figure 6-7a and figure 6-7b show the E_{ν}^{QE} and $\cos\theta$ distributions of the MiniBooNE excess, respectively, compared with the prediction from the neutrissimo model indicated by the black star in figure 6-6. The oscillation contribution from the 3 + 1 global fit without MiniBooNE is also shown.	210
6-8	The added time delay in MiniBooNE for a neutrissimo with the indicated parameters, as calculated in Ref. [27].	211
7-1	Schematic depiction of the Lujan TMRS Figure from Ref. [28].	231
7-2	Figure 7-3a, from Ref. [29], shows the energy distribution of π^+ decay-at-rest neutrinos from the Lujan beam dump source. Figure 7-2b, from Ref. [30], shows the timing distribution of particles produced in the Lujan beam dump source after traveling through the TMRS.	231
7-3	Figure 7-3a shows a schematic 3D rendering of the CCM200 detector. Figure 7-3b shows an image of the interior of the CCM200 detector.	233
7-4	Two of the veto PMT assemblies constructed at MIT, including the 1-inch PMT, base circuit board, and TPB-coated acrylic window.	234
7-5	Section 7.1 shows one of the veto PMTs across from the LED in the light-tight box. Section 7.1 shows the average response of the 20 veto PMTs to 1 V (top) and 2 V (bottom) LED pulses.	234

7-6	The timing distribution of photons from the Lujan source (solid black line) compared with that of neutrons measured in CCM120 (dashed red line). Figure from Ref. [30].	236
7-7	Figure 7-7a shows the integration in equation (7.1) as a function of λ_1 for $\lambda_2 = 700$ nm and $z = 1$. Figure 7-7b shows the Cherenkov cone angle $\cos \theta_C$ for an electron as a function of the photon wavelength and the electron kinetic energy.	239
7-8	Templates of the average number of p.e. detected in each CCM PMT within the first 8 ns of an electron event. Different templates are shown for electron kinetic energies of $T_{e^-} = 1$ MeV and $T_{e^-} = 5$ MeV, both with and without Cherenkov photons. Coated (uncoated) PMTs are indicated by the solid (dashed) circles. Grey PMTs indicate those which registered no hits within the first 8 ns across all simulations. The dimensions on each axis are in units of cm.	244
7-9	Example event displays showing the total number of p.e. detected in each CCM PMT within the first 8 ns of a single simulated electron event. Displays are shown for electron kinetic energies of $T_{e^-} = 1$ MeV and $T_{e^-} = 5$ MeV. Coated (uncoated) PMTs are indicated by the solid (dashed) circles. Grey PMTs indicate those which registered no hits in the first 8 ns of this specific simulation.	245
7-10	Distributions of the test statistic in equation (7.6) over all 10^4 simulations, considering either all photons or scintillation photons only. Distributions are shown for $T_{e^-} = 1$ MeV and $T_{e^-} = 5$ MeV. The vertical line indicates the lower bound requirement which can reject 99% of scintillation-only backgrounds.	245
7-11	Curves of the efficiency to retain events with Cherenkov light (“Ring efficiency”) v.s. the fraction of events without Cherenkov light that can be rejected (“No-Ring Rejection Factor”), generated by considering successively larger lower bounds on $\Delta \log \mathcal{L}$. Different curves are shown for $T_{e^-} \in \{1, 2, 3, 4, 5\}$ MeV.	246

7-12	Figure 7-12a shows a schematic depiction of the derivative-based pulse definition in the current CCM reconstruction. Figure 7-12b shows an example of this pulse finder in a data event, from Ref. [30].	249
7-13	Two example waveforms from CCM200 beam data. The regions identified by the derivative filter are indicated in red and the result of the fit to equation (7.7) is indicated by the purple curves.	249
7-14	Figure 7-14a shows an image of the six CosmicWatch pairs on top of the CCM detector. Section 7.2.2 shows a schematic diagram of the cosmic muon trigger in CCM200.	250
7-15	Figure 7-15a shows the summed waveform across all PMTs in CCM for a single example cosmic muon trigger. The delayed signal from a Michel electron can also be seen. Section 7.2.2 shows the difference in rise times between coated and uncoated PMT signals in the top and bottom halves of the barrel of the detector (labeled “sides-top” and “sides-bottom”, respectively), as described in the text.	250
7-16	Schematic depiction of prompt ν_μ from π^+ decay in the Lujan target upscattering to neutrinos within shielding along the path to CCM200 and decaying to photons in the detector. The pink circle represents the TMRS shown in figure 7-1.	255
7-17	Figure 7-17a shows the distribution of background and signal prediction in CCM200 for $m_N = 20.35$ MeV and $d_{\mu N} = 3 \times 10^{-7}$ GeV $^{-1}$, considering a background reduction factor of 10^{-3} compared to CCM120. Section 7.3 shows the background-subtracted plot, with a red band indicating the expected statistical uncertainty on the background. . .	256

7-18	Figure 7-18a shows the expected sensitivity of CCM200 to the neutrissimo model, where the blue band corresponds to a background reduction factor between 10^{-4} (“CCM200 Low Background”) and 10^{-2} (“CCM200 High Background”). The MiniBooNE E_ν^{QE} allowed region (pink) and existing constraints (grey) come from Ref [31]. Figure 7-18b shows the same plot, but considering $\mathcal{N} \rightarrow \nu_\tau \gamma$ decays with $d_{\tau\mathcal{N}} = d_{\mu\mathcal{N}} m_\tau / m_\mu$	256
7-19	The time delay of neutrissimo single photon decays with the CCM detector for the indicated mass and coupling, as calculated using <code>LeptonInjector</code> .257	
B-1	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	294
B-2	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	294
B-3	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	295
B-4	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	295
B-5	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	296
B-6	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	296
B-7	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	297
B-8	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	297
B-9	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	298
B-10	Top: pixel intensity; Bottom: <code>SparseSSNet</code> labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	298

B-11 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	299
B-12 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	299
B-13 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	300
B-14 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	300
B-15 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	301
B-16 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	301
B-17 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	302
B-18 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	302
B-19 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	303
B-20 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	303
B-21 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	304
B-22 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	304
B-23 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	305
B-24 Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex. . .	305

List of Tables

4.1	The definition of kinematic variables used throughout the two-body CCQE analysis.	142
4.2	The suite of variables used to isolate and analyze the $1e1p$ and $1\mu1p$ samples. Variables used in the BDT ensemble for each sample are specified. The “*” character indicates that the variable is calculated in the rest frame of the struck nucleon. The mathematical definitions of many of these variables appear in Table 4.1.	143
4.3	The specific used to define the $1\mu1p$ and $1e1p$ samples. For definitions of kinematic variables, see Table 4.1.	144
5.1	Breakdown of MC events in the “background” category of figures 5-3 and 5-4 over the range $200 < E_\nu < 1200$ MeV. The events are partitioned both by the interaction channel and the event topology.	152
5.2	Results from goodness-of-fit tests comparing observed $1e1p$ data to the H_0 and H_1 predictions, reported via the χ^2_{CNP} test statistic and the frequentist p -value. The top half of the table considers the nominal prediction and uncertainties before the $1\mu1p$ constraint described in section 5.1.3, while the bottom half considers the post-constraint prediction and uncertainties.	162
6.1	Relevant parameters of the ten most abundant nuclei in the Earth’s upper crust according to [32]	205

Chapter 1

Introduction

We begin with a brief primer on neutrinos, the surprises they have given physicists throughout recent history, and the mysteries that remain today. Readers already familiar with the mathematical details of massive neutrinos and the Standard Model may wish to read only section 1.1 and section 1.4 before continuing.

1.1 A Brief History of the Neutrino

The first indication of what would come to be known as the neutrino came from Wolfgang Pauli in 1930 [33]. Addressing the “radioactive ladies and gentlemen” of Tübingen, Germany, he appealed to the existence of new electrically neutral particles to save the law of energy conservation in nuclear beta decays. This idea was developed further by Enrico Fermi in 1934, who calculated the transition probability for β -decay with a neutrino in the final state [34]. Fermi’s theory represents the first study of the weak interaction—the only Standard Model gauge group under which neutrinos are charged.

As the name “weak interaction” suggests, neutrinos interact very feebly with particles in the Standard Model. Thus, it wasn’t until 1956 that the neutrino was observed in an experimental setting for the first time. A team of scientists from Los Alamos Scientific Laboratory, led by Frederick Reines and Clyde Cowan, detected a free neutrino from a nuclear reactor via the inverse beta decay interaction ($\bar{\nu}_e p \rightarrow e^+ n$) [35, 36].

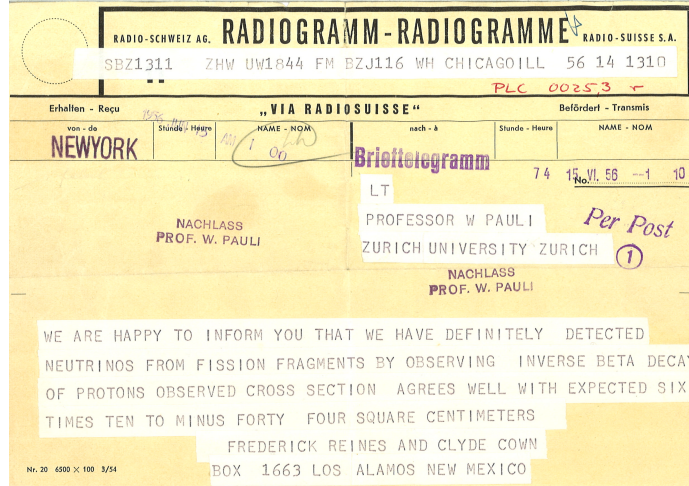


Figure 1-1: Telegram from Fred Reines and Clyde Cowan informing Wolfgang Pauli of their detection of neutrinos from a nuclear reactor.

Though it was not known at the time, they had detected electron antineutrinos ($\bar{\nu}_e$). Electron (anti)neutrinos represent one of the three weak-flavor eigenstates neutrinos can occupy in the Standard Model—specifically, the eigenstate that couples to the e^\pm charged leptons through the charged-current weak interaction. Upon confirmation of their discovery, Reines and Cowan sent the telegram shown in figure 1-1 to Pauli, alerting him of the definitive existence of the neutral particles he proposed in Tübingen.

Shortly after this, the phenomenon of neutrino oscillations—periodic transitions between different types of neutrinos—started to appear in the literature. In 1958, Bruno Pontecorvo discussed the possibility of mixing between right-handed antineutrinos $\bar{\nu}_R$ and “sterile” right-handed neutrinos ν_R , in analogy with $K^0-\bar{K}^0$ mixing observed in the quark sector [37]. A second possible source of neutrino oscillations came following the 1962 experimental discovery of a second neutrino weak-flavor eigenstate—the muon neutrino (ν_μ) [38]. After this, the notion of mixing between neutrino flavor and mass eigenstates was introduced by Ziro Maki, Masami Nakagawa, and Shoichi Sakata [39]. In a 1967 paper [40], Pontecorvo introduced the possibility of vacuum $\nu_e-\nu_\mu$ oscillations, even predicting a factor of two suppression in the total solar neutrino flux before such a deficit would actually be observed [41].

The aforementioned deficit, known as the “solar neutrino problem”, was estab-




Figure 1-2: The deficit of the observed solar ν_e flux compared with the theoretical expectation. The Homestake experiment is shown on the far left; follow-up solar neutrino measurements confirming the deficit are also shown, including the 2002 SNO result which brought forth a solution to the solar neutrino problem. Figure from Ref. [1].

lished in 1968 through a now-famous experiment at the Homestake Mine in South Dakota led by Raymond Davis [42]. Davis and his colleagues detected the capture of electron neutrinos from the sun on ^{37}Cl nuclei, allowing a measurement of the solar ν_e flux. Their result was about a factor of $\sim 1/3$ lower than the leading prediction from John Bahcall [43]. This is shown in figure 1-2, including confirmations of the deficit following the Homestake experiment. The solution was not a mistake in the experimental measurement or theoretical prediction, as physicists expected at the time; rather, it was a deficiency in our understanding of neutrinos. This was the first piece of the puzzle that would eventually lead to the discovery of neutrino oscillations and nonzero neutrino masses.

The next piece of the puzzle came from atmospheric neutrinos, i.e. neutrinos coming from the decay of mesons created from the interactions of primary cosmic rays in the atmosphere. Around the mid-1980s, two water Cherenkov detectors, IMB-3 [44] and Kamiokande [45], began to measure the interactions of atmospheric ν_μ and ν_e events (initially just as a background for their main physics goal, the search for nucleon decay). The ratio of $\nu_\mu : \nu_e$ interactions was found to be lower than the theoretical expectation by a factor of $\sim 2/3$ [46]. This was known as the “atmospheric neutrino anomaly”. The source of this anomaly was not clear at the time; it could

have been a deficit of muon neutrinos, an excess of electron neutrinos, or some of both. Systematic issues in the flux prediction or muon identification were also suggested [46]. It was far from clear that neutrino oscillations could be responsible for the observed deficit.

The solution to the atmospheric neutrino anomaly came from the Super-Kamiokande (SuperK) experiment [47]. SuperK was a much larger version of the Kamiokande detector, allowing the detection of higher energy muons (up to $E_\mu \sim 5$ GeV). SuperK also measured the up-down asymmetry of muon-like and electron-like events in their detector, $(N_{\text{up}} - N_{\text{down}})/(N_{\text{up}} + N_{\text{down}})$. Upward-going events have traveled a much longer distance than downward-going events before reaching the SuperK detector—thus positive detection of an asymmetry would be smoking-gun evidence for a baseline-dependent effect like neutrino oscillations. This is precisely what SuperK observed [2]. As shown in figure 1-3, an up-down asymmetry is observed in the muon-like channel, the magnitude of which increases with the observed muon momentum. Such behavior is consistent with muon neutrino oscillations to a third flavor eigenstate, ν_τ (the mathematical details of neutrino oscillations will be described in section 1.3). No such effect was observed in the electron-like channel. Thus, the atmospheric neutrino anomaly is a result of muon neutrino disappearance, specifically coming from $\nu_\mu \rightarrow \nu_\tau$ oscillations.

The solution to the solar neutrino problem came in 2002 from the Sudbury Neutrino Observatory (SNO) [48]. The SNO experiment used a heavy water Cherenkov detector, specifically relying on the use of deuterium target nuclei to be sensitive to three different neutrino interactions,

$$\begin{aligned}
 \nu_e + d &\rightarrow p + p + e^- & \text{(CC)}, \\
 \nu_x + d &\rightarrow p + n + \nu_x & \text{(NC)}, \\
 \nu_x + e^- &\rightarrow \nu_x + e^- & \text{(ES)}.
 \end{aligned}
 \tag{1.1}$$

Charged-current (CC), neutral-current (NC), and elastic scattering (ES) interactions were separated based on the visible energy and scattering angle of the final state

particles. NC events were further separated by tagging the 6.25 MeV photon released from neutron capture on deuterium. By measuring all three channels, SNO was able to measure the ^8B solar neutrino flux broken down into the ν_e and $\nu_{\mu,\tau}$ components. SNO's 2002 result showed that the missing neutrinos from the Homestake experiment were in fact showing up in the $\nu_{\mu,\tau}$ component [3]. Figure 1-4 shows the flux of each component as constrained by the measured CC, NC, and ES interaction rate. The flavor transitions here come not from vacuum oscillations but rather from matter-enhanced resonant behavior as neutrinos travel through the dense solar medium—a phenomenon known as the MikheyevSmirnovWolfenstein (MSW) effect [49, 50]. The MSW effect still, however, requires mixing between the neutrino flavor and mass eigenstate as well as non-zero squared differences between the mass eigenstates. It is worth noting here that the KamLAND reactor neutrino experiment was essential in determining the oscillation parameters which led to the SNO observation [51]. Thus, the SNO solution to the solar neutrino problem and the SuperK solution to the atmospheric neutrino anomaly were both evidence for the existence of neutrino oscillations and thus non-zero neutrino masses. The collaborations shared the 2015 Nobel Prize in physics for this discovery [52, 53].

Since SuperK and SNO, neutrino oscillations have been measured extensively by a global program of reactor, accelerator, atmospheric, and solar neutrino experiments. The mixing angle and mass-squared splittings of the three Standard Model neutrinos have been measured to few-percent-level precision in most cases [54–56]. There are a number of open questions in the standard three-neutrino mixing paradigm, including the ordering of the three mass eigenstates and the value of the charge-parity-violating complex phase δ_{CP} . Though preliminary results exist on both fronts [54–58], definitive answers to each will come from next-generation neutrino experiments, including Hyper-K [59], DUNE [60] and JUNO [61].




Figure 1-3: The up-down asymmetry measured in SuperK as a function of lepton momentum, separated into e -like and μ -like events as well as fully-contained (FC) and partially-contained (PC) events. The dashed line indicates the best fit to $\nu_\mu \rightarrow \nu_\tau$ oscillations. Figure from Ref. [2].




Figure 1-4: Measurement of the solar ^8B flux from the SNO collaboration, broken down into the ν_e and $\nu_{\mu,\tau}$ sub-components. Measurement of the CC, NC, and ES channels show up as slices in the two-dimensional flux parameter space. Figure from Ref. [3].

1.2 Neutrinos in the Standard Model

The arguments and notation presented in this section follow closely from chapter 2 of Ref. [62].

The interactions of the known fundamental particles of our Universe are described by a specific quantum field theory known as the Standard Model (SM). Above the electroweak scale, there are three gauge groups contained within the SM:

- $SU(3)_c$, which governs the gluon-mediated “strong interactions” of color-charged fields.
- $SU(2)_L$, one part of the “electro-weak interaction”, mediated by the W_μ^\pm and W_μ^0 vector bosons.
- $U(1)_Y$, the other part of the “electro-weak interaction”, mediated by the B_μ gauge boson.

After electro-weak symmetry breaking (EWSB) via the Higgs mechanism, the $SU(2)_L \times U(1)_Y$ subgroup breaks down to $U(1)_Q$, which describes the electromagnetic (EM) interactions of charged fields mediated by the A_μ gauge boson, also known as the photon.

Of the three fundamental interactions of the SM, neutrinos are only charged under the weak $SU(2)_L$ gauge group—they are singlets under the $SU(3)_c$ and $U(1)_Q$ gauge groups. Thus, neutrinos only appear in the electro-weak part of the SM Lagrangian, which is given by

$$\mathcal{L} = \frac{g}{\sqrt{2}}(J^\mu W_\mu^+ + J^{\mu\dagger} W_\mu^-) + \frac{g}{\cos\theta_W} K^\mu Z_\mu, \quad (1.2)$$

where $g = e/\cos\theta_W$ is the $SU(2)_L$ gauge coupling of W_μ and Higgs field, θ_W is the Weinberg angle describing the rotation that occurs during EWSB between the neutral parts of the $SU(2)_L$ and $U(1)_Y$ gauge boson fields, and W_μ^\pm (Z_μ) is the charged (neutral) piece of $SU(2)_L$ after EWSB. The currents coupled to W_μ^\pm and Z_μ bosons

are given by

$$\begin{aligned}
J^\mu &= \begin{pmatrix} \bar{u}^0 & \bar{c}^0 & \bar{t}^0 \end{pmatrix} \gamma^\mu P_L \begin{pmatrix} d^0 \\ s^0 \\ b^0 \end{pmatrix} + \begin{pmatrix} \bar{\nu}_e & \bar{\nu}_\mu & \bar{\nu}_\tau \end{pmatrix} \gamma^\mu P_L \begin{pmatrix} e \\ \mu \\ \tau \end{pmatrix} \\
K^\mu &= \sum_f \bar{f} \gamma^\mu [I_{3L} P_L - \sin^2 \theta_W Q_f] f \\
&= \sum_q [\epsilon_L(q) \bar{q} \gamma_\mu P_L q + \epsilon_R(q) \bar{q} \gamma_\mu P_R q] \\
&\quad + \frac{1}{2} \sum_{\alpha \in \{e, \mu, \tau\}} [\bar{\nu}_\alpha \gamma^\mu P_L \nu_\alpha + \bar{\ell}_\alpha \gamma_\mu (g_V^\alpha - \gamma_5 g_A^\alpha) \ell_\alpha],
\end{aligned} \tag{1.3}$$

where $P_R(P_L) = (1 \pm \gamma^5)/2$ is the projection operator onto the right-handed (left-handed) chiral state, and the subscript 0 on the quark fields indicates that these are the weak flavor eigenstates rather than the mass eigenstates. The first generation coupling constants in K^μ , which derive from the specified EM charge and $SU(2)_L$ representation of each field, are given by

$$\begin{aligned}
\epsilon_L(u) &= \frac{1}{2} - \frac{2}{3} \sin^2 \theta_W & \epsilon_R(u) &= -\frac{2}{3} \sin^2 \theta_W \\
\epsilon_L(d) &= -\frac{1}{2} + \frac{1}{3} \sin^2 \theta_W & \epsilon_R(d) &= \frac{1}{3} \sin^2 \theta_W \\
g_V^e &= -\frac{1}{2} + 2 \sin^2 \theta_W & g_A^e &= -\frac{1}{2}.
\end{aligned} \tag{1.4}$$

The Lagrangian in equation (1.2) can be used to calculate cross sections for the various SM interactions of the neutrino. The first term describes the charged-current interactions of neutrinos such as nuclear beta decay, while the second term describes neutral current interactions such as $\nu_\mu e^-$ elastic scattering. At energy scales below the electro-weak scale, one can integrate out the W_μ and Z_μ gauge bosons and describe interactions in terms of the dimensional Fermi constant

$$G_F = \frac{g^2}{4\sqrt{2}M_W^2} = 1.166 \times 10^{-5} \text{ GeV}^{-2}. \tag{1.5}$$

The low-energy Lagrangian describing 4-fermion interactions can be derived from

equation (1.2) as

$$\mathcal{L}_{Af} = \frac{-4G_F}{\sqrt{2}} [J_\mu J^{\mu\dagger} + K_\mu K^\mu]. \quad (1.6)$$

As an example, we consider low-energy neutrino electron elastic scattering (ES) ($\nu e^- \rightarrow \nu e^-$). This is a purely leptonic process and is therefore relatively clean; specifically, ES models do not need to account for the complex dynamics of the nuclear medium. The Feynman diagrams for the contributing interactions are shown in figure 1-5. Both the charged-current (CC) and neutral-current (NC) diagrams contribute to $\nu_e e^-$ scattering, while only the NC diagram contributes to $\nu_{\mu,\tau} e^-$ scattering. Using the Feynman rules associated with equation (1.6), one can calculate the cross sections to be [62]

$$\begin{aligned} \sigma_{\nu_e e^- \rightarrow \nu_e e^-}(E_\nu) &= \frac{G_F^2 m_e E_\nu}{2\pi} \left[(2 \sin^2 \theta_W + 1)^2 + \frac{4}{3} \sin^4 \theta_W \right] \\ &\approx 0.9 \times 10^{-43} \left(\frac{E_\nu}{10 \text{ MeV}} \right) \text{cm}^2 \\ \sigma_{\nu_{\mu,\tau} e^- \rightarrow \nu_{\mu,\tau} e^-}(E_\nu) &= \frac{G_F^2 m_e E_\nu}{2\pi} \left[(2 \sin^2 \theta_W - 1)^2 + \frac{4}{3} \sin^4 \theta_W \right] \\ &\approx 0.15 \times 10^{-43} \left(\frac{E_\nu}{10 \text{ MeV}} \right) \text{cm}^2, \end{aligned} \quad (1.7)$$

which is valid for $E_\nu \gg m_e$. Similarly, one can calculate the cross section for antineutrino electron ES ($\bar{\nu} e^- \rightarrow \bar{\nu} e^-$). The diagrams contributing for this process are shown in figure 1-6, and the cross section is given by [62]

$$\begin{aligned} \sigma_{\bar{\nu}_e e^- \rightarrow \bar{\nu}_e e^-}(E_\nu) &= \frac{G_F^2 m_e E_\nu}{2\pi} \left[\frac{1}{3} (2 \sin^2 \theta_W + 1)^2 + 4 \sin^4 \theta_W \right] \\ &\approx 0.378 \times 10^{-43} \left(\frac{E_\nu}{10 \text{ MeV}} \right) \text{cm}^2 \\ \sigma_{\bar{\nu}_{\mu,\tau} e^- \rightarrow \bar{\nu}_{\mu,\tau} e^-}(E_\nu) &= \frac{G_F^2 m_e E_\nu}{2\pi} \left[\frac{1}{3} (2 \sin^2 \theta_W - 1)^2 + 4 \sin^4 \theta_W \right] \\ &\approx 0.14 \times 10^{-43} \left(\frac{E_\nu}{10 \text{ MeV}} \right) \text{cm}^2. \end{aligned} \quad (1.8)$$

We now turn to the interaction at the core of this thesis: neutrino-nucleon charged-current quasi-elastic (CCQE) scattering. The relevant Feynman diagrams for this

process are shown in figure 1-7. Unlike ES, models of CCQE do need to account for the nuclear environment surrounding the target nucleon. As the final state nucleon travels through the nuclear medium, it may scatter off of other nucleons and/or produce additional mesons through a process known as final state interactions (FSIs). As shown in figure 1-8, CCQE is dominant for $E_\nu \lesssim 1$ GeV. Above this energy, nucleon resonance processes start to take over, in which Delta resonances decay to final state mesons. In the regime $E_\nu \gtrsim 10$ GeV, neutrinos start to undergo deep inelastic scattering (DIS) off of the constituent quarks within the nucleon.

In order to calculate the CCQE cross section, one considers a theory containing nucleon degrees of freedom. The original calculation for free nucleons (i.e., not bound within a nucleus) was carried out by Llewellyn-Smith in 1972; the differential cross section as a function of the squared four-momentum transfer Q^2 is given by [4, 63]

$$\frac{d\sigma}{dQ^2} = \frac{G_F^2 M^2 |V_{ud}|^2}{8\pi E_\nu^2} \left[A \pm \frac{s-u}{M^2} B + \frac{(s-u)^2}{M^4} C \right], \quad (1.9)$$

where $+(-)$ refers to (anti)neutrino scattering, M is the nucleon mass, m is the lepton mass, $(s-u) = 4ME_\nu - Q^2 - m^2$, and A , B , and C are functions of the vector, axial-vector, and pseudoscalar form factors of the nucleon (see equations 58, 59, and 60 of Ref. [4] for complete expressions). These form factors describe the composite nature of nucleons under interactions with different Lorentz structures.

For $E_\nu \ll M$, the ν_e CCQE cross section is approximately [62]

$$\begin{aligned} \sigma_{\nu_e n \rightarrow e^- p}(E_\nu) &\approx \frac{G_F^2 E_\nu^2}{\pi} (g_V^2 + 3g_A^2) \\ &\approx 9.75 \times 10^{-42} \left[\frac{E_\nu}{10 \text{ MeV}} \right]^2 \text{ cm}^2. \end{aligned} \quad (1.10)$$

In the regime $E_\nu \gtrsim 1$ GeV, the ν_e and ν_μ CCQE cross sections are no longer suppressed by threshold effects and are thus the same, approximately 10^{-38} cm² [4]. This cross section is significantly larger than the elastic scattering and lower energy ν_e CCQE cross sections and is the dominant neutrino interaction for many accelerator-based neutrino experiments, including two at the heart of this thesis: MiniBooNE and




Figure 1-5: Diagrams contributing to νe^- elastic scattering




Figure 1-6: Diagrams contributing to $\bar{\nu} e^-$ elastic scattering

MicroBooNE. Finally, we note that the cross section for antineutrino CCQE tends to be smaller; this will be important in chapter 6.

1.3 Massive Neutrinos

The arguments and notation presented in this section follow closely from section 2.5, chapter 4, and chapter 5 of Ref. [62] as well as chapter 11 of Ref. [64].

Neutrinos are massless in the SM. To see this, we will exhaust the two possible forms for a neutrino mass term in the SM Lagrangian: Dirac and Majorana. These refer to the two possible fermionic spinor representations in which neutrinos can be found. Dirac spinors in general have four complex components, or degrees of freedom, while Majorana spinors have only two. The critical question is whether the right-handed chiral projection of the neutrino field, ν_R , is the same as $\bar{\nu}_R$, the right-handed chiral projection of the antineutrino field (Majorana case), or if it is a new degree of




Figure 1-7: Diagrams contributing to neutrino-nucleon charged-current quasielastic scattering




Figure 1-8: CC inclusive neutrino and antineutrino nucleon scattering cross sections as a function of neutrino energy. Figure from Ref. [4].

freedom (Dirac case).

The definition of a free Dirac fermion field is

$$\psi(x) = \int \frac{d^3p}{\sqrt{(2\pi)^3 2E_p}} \sum_{s=\pm\frac{1}{2}} \left(f_s(\mathbf{p}) u_s(\mathbf{p}) e^{-i\mathbf{p}\cdot\mathbf{x}} + \bar{f}_s^\dagger(\mathbf{p}) v_s(\mathbf{p}) e^{i\mathbf{p}\cdot\mathbf{x}} \right), \quad (1.11)$$

where $f_s(\mathbf{p})$ annihilates a single particle of momentum \mathbf{p} while \bar{f}_s^\dagger creates the corresponding antiparticle state, and $u_s(\mathbf{p})$ and $v_s(\mathbf{p})$ are spinors with positive and negative energy, respectively, satisfying the Dirac equations

$$\begin{aligned} (\gamma^\mu p_\mu - m) u_s(\mathbf{p}) &= 0 \\ (\gamma^\mu p_\mu + m) v_s(\mathbf{p}) &= 0, \end{aligned} \quad (1.12)$$

where γ^μ are a set of Lorentz-indexed matrices satisfying $\{\gamma^\mu, \gamma^\nu\} = 2g^{\mu\nu}$. There are many possible representations for the γ -matrices. We consider the Weyl basis, in which [64]

$$\gamma_\mu = \begin{pmatrix} 0 & \sigma_\mu \\ \bar{\sigma}_\mu & 0 \end{pmatrix}, \quad (1.13)$$

where $\sigma_\mu = (\mathbb{1}, \vec{\sigma})$, $\bar{\sigma}_\mu = (\mathbb{1}, -\vec{\sigma})$, and $\vec{\sigma} = (\sigma_1, \sigma_2, \sigma_3)$ are the Pauli matrices. This representation is convenient for understanding the different chiral components of the Dirac spinor ψ . The Lorentz generators $S^{\mu\nu} \equiv \frac{i}{4}[\gamma^\mu, \gamma^\nu]$ become block diagonal, such that we can write the Dirac spinor of equation (1.11) as a doublet of two-component Weyl spinors with a different chiral nature,

$$\psi = \begin{pmatrix} \psi_L \\ \psi_R \end{pmatrix}, \quad (1.14)$$

which transform under different irreducible representations of the Lorentz group [64]. Here ψ_L and ψ_R refer to the left-handed and right-handed Weyl spinor, respectively. We can isolate the different chiral components of the Dirac spinor using the matrix

$\gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3$, which takes the form

$$\gamma^5 = \begin{pmatrix} -\mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix} \quad (1.15)$$

in the Weyl basis. We can define projection operators $P_L = \frac{1}{2}(1 - \gamma^5)$ and $P_R = \frac{1}{2}(1 + \gamma^5)$ such that $P_L\psi = \psi_L$ and $P_R\psi = \psi_R$. It is worth noting that while the behavior of these projection operators is especially clear in the Weyl representation, they will isolate the chiral components of ψ in any representation of γ^μ .

Dirac mass terms couple left-handed and right-handed chiral states. To see this, consider the Dirac equation in the Weyl basis, which takes the form [64]

$$(\gamma^\mu p_\mu - m)\psi = \begin{pmatrix} -m & \sigma^\mu p_\mu \\ \bar{\sigma}^\mu p_\mu & -m \end{pmatrix} \begin{pmatrix} \psi_L \\ \psi_R \end{pmatrix} = 0. \quad (1.16)$$

It is evident that this matrix equation mixes the left-handed and right-handed components of ψ . Dirac mass terms take the form $m\psi_L^\dagger\psi_R$ and $m\psi_R^\dagger\psi_L$, thus requiring both chiral components. After EWSB, the non-neutrino fermions in the SM acquire a Dirac mass from their interactions with the Higgs field. Neutrinos, however, do not have a right-handed chiral state in the SM; therefore, the SM cannot include a Dirac mass term for neutrinos.

Now we turn to the Majorana mass term. The expression for a Majorana field is the same as equation (1.11), subject to a condition relating particles and antiparticles. We see that the expression for $\psi^*(x)$ would involve $f_s^\dagger(\mathbf{p})$, which creates a particle state, and $\bar{f}_s(\mathbf{p})$, which annihilates an antiparticle state. It turns out the relationship $\psi(x) = \psi^*(x)$ is not Lorentz invariant [62]. To remedy this, we must define the conjugate Dirac field

$$\psi^C(x) \equiv \gamma_0 C \psi^*(x), \quad (1.17)$$

where the representation-dependent conjugation matrix C is defined by the equation

$$\begin{aligned}\gamma_0 C \sigma_{\mu\nu}^* &= -\sigma_{\mu\nu} \gamma_0 C, \\ \sigma_{\mu\nu} &\equiv \frac{i}{2} [\gamma_\mu, \gamma_\nu].\end{aligned}\tag{1.18}$$

In the Weyl representation, for example, $C = i\gamma_2\gamma_0$. This requirement for C ensures that $\psi^C(x)$ transforms in the same way as $\psi(x)$ under the Lorentz group [62]. The Lorentz-invariant Majorana condition specifically requires

$$\psi(x) = e^{i\theta} \psi^C(x),\tag{1.19}$$

where θ is an arbitrary phase, which we can take to be $\theta = 0$. It is important to note that this condition can only be satisfied for fields that carry no additive quantum numbers [64].

In the Weyl basis, equation (1.19) relates the left-handed and right-handed components of $\psi(x)$ such that [64]

$$\psi(x) = \begin{pmatrix} \psi_L \\ i\sigma_2 \psi_L^* \end{pmatrix},\tag{1.20}$$

where the number of degrees of freedom has been reduced from four to two. Since $i\sigma_2 \psi_L^*$ transforms like a right-handed spinor, we can now write mass terms of the form $im\psi_L^\dagger \sigma_2 \psi_L^*$ and $im\psi_L^\dagger \sigma_2 \psi_L^*$. These are Majorana mass terms. Note that they couple the same chiral component of the fermion.

The impossibility of a neutrino Majorana mass term is a bit more nuanced. Majorana mass terms for neutrinos in the SM contain the bi-linear expression $\nu_L^T \sigma_2 \nu_L$. However, ν_L belongs to an $SU(2)_L$ doublet in the SM, thus this Majorana mass term transforms as a triplet under $SU(2)_L$. It also breaks lepton number by two units, hence it also violates baryon minus lepton number ($B - L$), which is conserved to all orders of the SM gauge couplings [62]. Therefore, neutrinos also cannot have a Majorana mass term in the SM.

Despite these arguments, neutrino oscillations have given physicists definitive evidence that at least two of the three SM neutrino masses are nonzero (as discussed in section 1.1). This requires the presence of physics beyond the Standard Model (BSM). The minimal extension of the SM which can accommodate nonzero neutrino masses introduces additional right-handed neutrino states N_R [62, 64]. These fields, which are singlets under the SM gauge group, can generate both Dirac and Majorana mass terms for neutrinos. The most general expression for the neutrino mass Lagrangian is then

$$-\mathcal{L}_{\text{mass}} = \frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \overline{N_L^C} \end{pmatrix} \begin{pmatrix} 0 & M \\ M^T & B \end{pmatrix} \begin{pmatrix} \nu_R^C \\ N_R \end{pmatrix} + \text{h.c.}, \quad (1.21)$$

where M and B are the Dirac and Majorana mass matrices of the neutrino sector, respectively, and ν_L and N_R are column vectors containing the left-handed and right-handed projections of each neutrino generation.

In order to obtain the mass eigenstates of this theory, one must diagonalize the mass matrix in equation (1.21). If we assume one generation of neutrinos, the eigenvalues of this mass matrix are

$$m_{1,2} = \frac{1}{2} (\sqrt{B^2 + 4M^2} \mp B). \quad (1.22)$$

In the limit $B \gg M$, the eigenvalues are approximately given by

$$m_1 \approx \frac{M^2}{B}, \quad m_2 \approx B. \quad (1.23)$$

This is the famous “seesaw mechanism” for neutrino mass generation [65]. One can see that if B is at roughly the GUT scale (10^{16} GeV) and M is at roughly the electro-weak scale (100 GeV), we see that $m_1 < 1$ eV. This is the right order-of-magnitude regime predicted by neutrino oscillation data and is consistent with existing upper bounds on the neutrino mass from KATRIN [66]. Thus, this model is an elegant explanation of the observed neutrino oscillation phenomenon, though experimental confirmation of right-handed neutrino fields at the GUT scale is probably not feasible for quite a long time.

While we do not know the mechanism through which neutrinos acquire mass, it is relevant to ask whether the resulting mass terms are Dirac or Majorana in nature. An extensive worldwide experimental program is currently underway to answer this question by searching for neutrino-less double beta decay, a rare decay process in which a nucleus undergoes two simultaneous beta decays without emitting any neutrinos in the final state [67–69]. A positive observation would imply that neutrinos are Majorana.

As discussed in section 1.1, perhaps the most famous consequence of massive neutrinos is the phenomenon of neutrino oscillations [37, 39, 40]. This arises because the three weak flavor eigenstates ν_α are not aligned with the three mass eigenstates ν_i . The two bases are related by the unitary Pontecorvo–Maki–Nakagawa–Sakata (PMNS) mixing matrix $U_{\alpha i}$,

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \quad (1.24)$$

As seen in equation (1.2), neutrinos interact in the weak flavor eigenstates ν_α . Thus, a neutrino produced alongside a charged anti-lepton $\bar{\ell}$ is in the state

$$|\nu(t=0)\rangle = |\nu_\ell\rangle = \sum_{i \in \{1,2,3\}} U_{\ell i} |\nu_i\rangle. \quad (1.25)$$

Neutrinos propagate, however, in their mass eigenstates. Each mass eigenstate ν_i is associated with a mass m_i and four-momentum $(p_i)_\mu = (E_i, \vec{p}_i)$ satisfying the on-shell requirement $(p_i)^2 = m_i^2$. Thus, after a time t , the neutrino will be in the state

$$|\nu(t)\rangle = \sum_i e^{-ip_i \cdot x} U_{\ell i} |\nu_i\rangle. \quad (1.26)$$

The overlap with a different weak flavor eigenstate $\nu_{\ell'} \neq \nu_\ell$ is non-trivial, given

by the expression

$$\begin{aligned}
\langle \nu_{\ell'} | \nu(t) \rangle &= \sum_{i,j} \langle \nu_j | U_{j\ell'}^\dagger e^{-ip_i \cdot x} U_{\ell i} | \nu_i \rangle \\
&= \sum_i e^{-ip_i \cdot x} U_{\ell i} U_{\ell' i}^*,
\end{aligned} \tag{1.27}$$

where we have invoked the orthonormality of the mass basis in the last line. The probability of finding a neutrino in flavor eigenstate $\nu_{\ell'}$ given an initial ν_ℓ state is then

$$\begin{aligned}
P_{\nu_\ell \rightarrow \nu_{\ell'}}(t) &= |\langle \nu_{\ell'} | \nu(t) \rangle|^2 \\
&= \sum_{i,j} |U_{\ell i} U_{\ell' i}^* U_{\ell j}^* U_{\ell' j}| e^{-i(p_i - p_j) \cdot x + i\phi_{\ell\ell'ij}}.
\end{aligned} \tag{1.28}$$

where $\phi_{\ell\ell'ij} \equiv \arg(U_{\ell i} U_{\ell' i}^* U_{\ell j}^* U_{\ell' j})$.

We now make a simplifying assumption, in which all neutrino mass eigenstates propagate with the same momentum, i.e. $\vec{p}_i = \vec{p}_j \equiv \vec{p} \forall i, j$. This treatment is not necessarily physical. However, for the parameters relevant to most laboratory neutrino experiments, it leads to the same result as the correct but complicated full treatment of the quantum mechanical neutrino wave packet [70]. Given this assumption along with the approximation that $m_i \ll p_i$ (which should hold for all existing and near-future experiments), we can show

$$\begin{aligned}
(p_i - p_j) \cdot x &= (E_i - E_j)t \\
&= \left(\sqrt{\vec{p}^2 + m_i^2} - \sqrt{\vec{p}^2 + m_j^2} \right) t \\
&\approx \frac{\Delta m_{ij}^2 t}{2|\vec{p}|},
\end{aligned} \tag{1.29}$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$. Working in natural units ($c = \hbar = 1$), we note that ultra-relativistic neutrinos satisfy $|\vec{p}| \approx E$ and $t \approx L$, where L is the distance traveled by

the neutrino. Taking only the real part of the exponential in equation (1.28), we have

$$P_{\nu_\ell \rightarrow \nu_{\ell'}}(t) = \sum_{i,j} |U_{\ell i} U_{\ell' i}^* U_{\ell j}^* U_{\ell' j}| \cos \left(\frac{\Delta m_{ij}^2 L}{2E} - \phi_{\ell \ell' ij} \right). \quad (1.30)$$

If we consider a two-neutrino paradigm, the unitary mixing matrix is real and can be parameterized by a single “mixing angle” θ ,

$$U \equiv \begin{pmatrix} U_{\ell 1} & U_{\ell 2} \\ U_{\ell' 1} & U_{\ell' 2} \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}. \quad (1.31)$$

In this scenario, summing over the two mass eigenstates as in equation (1.30) gives

$$P_{\nu_\ell \rightarrow \nu_{\ell'}}(t) = \sin^2 2\theta \sin^2 \left(\frac{\Delta m^2 L}{4E} \right). \quad (1.32)$$

The extension to the standard three neutrino paradigm can be found in any text on neutrino oscillations. We quote the result here. Three mass eigenstates lead to two independent mass-squared splittings, Δm_{12}^2 and Δm_{23}^2 . The mixing matrix in equation (1.24) can be parameterized by three real mixing angles θ_{ij} and one complex CP-violating phase δ ,

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13} e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13} e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.33)$$

where $c_{ij} \equiv \cos \theta_{ij}$ and $s_{ij} \equiv \sin \theta_{ij}$. The three mixing angles (θ_{12} , θ_{13} , θ_{23}) and two relevant mass squared splittings Δm_{12}^2 and $|\Delta m_{23}^2|$ have been measured to a precision of $\mathcal{O}(1\% - 10\%)$ over the past two decades [54–56]. An extensive experimental program is planned to measure δ to similar precision, as well as the neutrino hierarchy (i.e., the sign of Δm_{23}^2) and the octant of θ_{23} [71].

1.4 Anomalies in the Neutrino Sector

Despite the success of the three-neutrino mixing paradigm, several anomalous results have appeared. Perhaps the most famous of these is the excess of $\bar{\nu}_e$ candidate events observed by the Liquid Scintillator Neutrino Detector (LSND) experiment [5]. LSND took data at Los Alamos Meson Physics Facility (LAMPF) from 1993-1998, observing neutrino interactions from a high-intensity decay-at-rest (DAR) source. The LSND detector was a 167-ton cylindrical tank of mineral oil that collected scintillation and Cherenkov light produced in neutrino interactions. The LAMPF accelerator provided a ~ 1 mA beam of 798 MeV protons, which were then focused on a water or high-Z target. This process created a large number of pions, which then decayed to produce neutrinos. Most π^- came to rest and were captured by nuclei in and around the target, and the $\pi^+ \rightarrow \nu_e e^+$ decay chain is helicity-suppressed due to the interplay between angular momentum conservation and the left-chiral nature of the weak interaction. Thus the dominant neutrino production process was $\pi^+ \rightarrow \nu_\mu (\mu^+ \rightarrow \bar{\nu}_\mu \nu_e e^+)$.

LSND looked specifically for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ conversion using $\bar{\nu}_\mu$ from μ^+ DAR. The $\bar{\nu}_e$ events were observed via the inverse beta decay (IBD) process. This is a very clean channel, as one can require a coincidence between the initial positron emission and the subsequent neutron capture on hydrogen, which releases a characteristic 2.2 MeV photon. The intrinsic $\bar{\nu}_e$ flux, coming predominately from π^- decay-in-flight (DIF), was suppressed compared to intrinsic $\bar{\nu}_\mu$ by a factor of $\sim 8 \times 10^{-4}$. Any significant excess of $\bar{\nu}_e$ would be evidence for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations. This is exactly what LSND observed, as shown in figure 1-9. However, the neutrino energies $\mathcal{O}(30$ MeV) and baselines ($\mathcal{O}(30$ m) required a mass-squared-splitting of $\Delta m^2 \sim 1$ eV². This is larger than the measured values of Δm_{12}^2 and $|\Delta m_{23}^2|$ by at least three orders of magnitude—therefore, the LSND result cannot be explained by the standard three neutrino oscillation paradigm. One must introduce a fourth neutrino to the SM neutrinos in order to facilitate such oscillations. Measurements of the invisible width of the Z boson forbid this neutrino from coupling to the weak force in the same way as the three SM neutrinos [72]. Thus, this fourth neutrino is typically referred to as a




Figure 1-9: The LSND excess of $\bar{\nu}_e$ events on top of the predicted SM background (green and red regions). The blue region indicates the best fit to $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations via a sterile neutrino state. Figure from Ref. [5]

“sterile neutrino” (ν_s). The sterile neutrino paradigm will be introduced in more detail in section 1.4 and discussed thoroughly throughout this thesis. The LSND anomaly is currently under direct investigation by the follow-up JSNS² experiment [73, 74], which will use a gadolinium-loaded liquid scintillator detector [75] to measure IBD interactions at the J-PARC Materials and Life Science Experimental Facility.

The Mini Booster Neutrino Experiment (MiniBooNE) was designed to follow up on the LSND anomaly [76]. MiniBooNE took data at Fermilab’s Booster Neutrino Beam (BNB) from 2002-2019, observing the interactions of neutrinos with energy $E \sim 500$ MeV in an 800-metric-ton mineral oil (CH_2) detector [15]. The Fermilab Booster accelerates protons to a kinetic energy of 8 GeV, at which point they collide with the beryllium target of the BNB. This produces a cascade of mesons, predominately pions. The charged mesons are focused using a magnetic horn and decay in a 50 m decay pipe; in the nominal “neutrino mode”, the magnetic field is generated to create a flux of mostly muon neutrinos from π^+ decay-in-flight [14]. The current in the magnetic horns can be reversed to instead focus π^- along the beamline, thus creating a beam

of mostly muon antineutrinos—this is referred to as “antineutrino mode”. MiniBooNE was situated at a baseline of $L \sim 500$ m from the BNB target, resulting in a similar characteristic L/E as that of LSND, ≈ 1 m/MeV. By equation (1.30), this means MiniBooNE would also be sensitive to oscillations at $\Delta m^2 \sim 1$ eV².

In 2007, MiniBooNE began to report results from their flagship analysis: the search for an excess of ν_e events in the BNB [76]. MiniBooNE relied primarily on the reconstruction of Cherenkov light from charged final state particles to identify neutrino interactions. Thus, ν_e CC interactions would show up as a “fuzzy” Cherenkov ring due to multiple scattering of the electron as well as the induced EM shower [77]. These fuzzy Cherenkov ring events are hereafter referred to as “electron-like” events. Starting with the initial results [76,78], MiniBooNE has consistently observed an excess of electron-like events above their expected SM background, the significance of which has grown over the 17-year data-taking campaign of the experiment [10]. Figure 2-5 shows the 4.8σ MiniBooNE electron-like excess considering the total neutrino mode dataset, corresponding to 18.75×10^{20} protons-on-target (POT) [10]. A similar excess was observed in the antineutrino mode dataset [79]. The as-yet-unexplained MiniBooNE excess represents one of the most significant disagreements with the SM to date.

Though the origin of the MiniBooNE excess remains unknown, neutrino physicists have converged on a number of potential explanations. The most famous explanation involves sterile neutrino-driven $\nu_\mu \rightarrow \nu_e$ oscillations consistent with the LSND result ($\Delta m^2 \sim 1$ eV²). While this model can explain at least some of the MiniBooNE excess, the excess in the lowest energy region ($E_\nu \lesssim 400$ MeV) sits above even the best-fit sterile neutrino solution. Due to the Cherenkov nature of the detector, electrons and photons are essentially indistinguishable—both seed EM showers which appear as fuzzy Cherenkov rings. Thus, the MiniBooNE excess could also come from a mismodeled photon background. Though not the subject of this thesis, there have been extensive experimental and theoretical efforts, both within and outside of the MiniBooNE collaboration, to validate the MiniBooNE SM photon background prediction [10,80–82]. One can also consider BSM sources of electron-like events in




Figure 1-10: The MiniBooNE electron-like channel data and SM background prediction for the entire neutrino mode dataset, as a function of the reconstructed neutrino energy.

MiniBooNE. Typical models introduce additional sources of photons and/or e^+e^- events in MiniBooNE through couplings to new dark sector particles. Resolution of the LSND and MiniBooNE anomalies, often referred to as the short baseline (SBL) anomalies, is a major goal within the particle physics community [83]. This thesis specifically investigates the MiniBooNE anomaly in further detail, covering both experimental and phenomenological studies into the origin of the excess.

We now briefly touch on two additional classes of anomalies that have surfaced over the years: the reactor antineutrino anomaly (RAA) and the gallium anomaly. The RAA [8] is a $\sim 5\%$ deficit in the total $\bar{\nu}_e$ rate observed from nuclear reactors compared to the theoretical expectation from the Huber-Mueller (HM) model [84, 85]. The HM model combines results using the summation method (summing the contributions of all beta-decay branches in the reactor) and the conversion method (relying on older measurements of the $\bar{\nu}_e$ flux from the different fissionable isotopes in the reactor). The data contributing to the RAA mostly come from reactor neutrino experiments operating at baselines short enough that the effects of SM neutrino oscillations are negligible. One can interpret the RAA as due to $\bar{\nu}_e$ disappearance via oscillations involving a sterile neutrino. Coincidentally, due to the relevant neutrino energies and baselines, such a solution requires $\Delta m^2 \gtrsim 1 \text{ eV}^2$, similar to the LSND and MiniBooNE




Figure 1-11: Data contributing to the reactor antineutrino anomaly, indicating the $\sim 5\%$ flux deficit observed by short-baseline reactor neutrino experiments. The red line indicates the prediction incorporating SM neutrino oscillations only, while the blue line shows an example prediction including a sterile neutrino. Figure from Ref. [6].

solution [6]. Figure 1-11 shows an overview of the RAA circa 2012, including the suite of short baseline reactor experiments which observe a deficit with respect to the HM model with SM neutrino oscillations (red line), as well as an example sterile neutrino solution to the RAA (blue line). Recently, the reactor $\bar{\nu}_e$ flux calculation has been revisited by various groups, each of which improves upon some aspect of the summation or conversion method used in the HM flux model [86–89]. The significance of the RAA either diminishes or disappears in some of these models; however, these improved models have difficulty removing the RAA while also explaining the “5-MeV bump” observed by most short baseline reactor experiments with respect to the HM model [89]. Thus, while the RAA story is quickly evolving, our understanding of reactor neutrino fluxes is far from clear.

The gallium anomaly refers to a series of gallium-based detectors that have observed a deficit of ν_e capture events on ^{71}Ga with respect to the theoretical expectation. The original harbingers of the anomaly, SAGE [90] and GALLEX [91], were designed to measure solar neutrinos using the $^{71}\text{Ga}\nu_e \rightarrow ^{71}\text{Ge}e^-$ capture process. Each detector was calibrated using electron capture ν_e sources, including ^{51}Cr and ^{37}Ar . Combining all available calibration data across both experiments, the observed ^{71}Ge production rate was lower than the expectation by a factor of 0.87 ± 0.05 [90]. Though the statistical significance of the anomaly was only modest ($2 - 3\sigma$), the community was already beginning to interpret the anomaly as $\nu_e \rightarrow \nu_s$ transitions via an eV-scale sterile neutrino [92]. A follow-up experiment to the SAGE and GALLEX anomaly, BEST [9], released their first results in 2021. BEST placed a 3.414 MCI




Figure 1-12: Data contributing to the gallium anomaly, indicating the $\sim 20\%$ deficit in the ^{71}Ge production rate observed by SAGE, GALLEX, and BEST. Figure from Ref. [7].

^{51}Cr ν_e source at the center of two nested ^{71}Ga volumes, each with a different average distance from the source. The ratio of observed to the predicted ^{71}Ge production rate was $R_{in} = 0.79 \pm 0.05$ ($R_{out} = 0.77 \pm 0.05$) for the inner (outer) volume, thus reaffirming the gallium anomaly [9]. No evidence for a difference in the deficit between the inner and outer volumes was observed, which would have been a smoking gun signature of a baseline-dependent effect like $\nu_e \rightarrow \nu_s$ oscillations. However, the statistical significance of the gallium anomaly is now much stronger; the combined SAGE, GALLEX, and BEST results give evidence for a deficit at the 5.0σ level [7]. The datasets contributing to this anomaly are summarized in figure 1-12.

As alluded to above, the most common BSM interpretation of the SBL, reactor antineutrino, and gallium anomalies is the “3+1 model”, which involves the addition of a new neutrino state—the sterile neutrino—at the eV scale. The sterile neutrino introduces a fourth weak interaction eigenstate ν_s and mass eigenstate ν_4 to the

standard three-neutrino mixing paradigm. Thus, equation (1.24) becomes

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \\ \nu_s \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} & U_{e4} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} & U_{\mu4} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} & U_{\tau4} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \nu_4 \end{pmatrix}. \quad (1.34)$$

As we are interested in an eV-scale sterile neutrino, the mass-squared splittings between the three active neutrinos are smaller by at least 2-3 orders of magnitude compared to their mass-squared splittings with the fourth mass eigenstate. This means that the active neutrino mass splittings are negligible for short-baseline experiments, i.e. those in which the argument of the second \sin^2 term in equation (1.32) is small. Experiments contributing to the aforementioned anomalies all satisfy this condition. Thus, when considering sterile neutrino explanations for these anomalies, we can make the approximation

$$\Delta m_{41}^2 \approx \Delta m_{42}^2 \approx \Delta m_{43}^2 \equiv \Delta m^2, \quad (1.35)$$

where we hereafter use Δm^2 to refer to the mass-squared splitting of the fourth mass eigenstate. This approximation holds regardless of the hierarchy of SM neutrino mass eigenstates.

The experiments discussed in this thesis are sensitive only to $\vec{\nu}_e$ and $\vec{\nu}_\mu$ interactions. The sterile neutrino can facilitate short-baseline oscillations between these flavor states; the oscillation probability expressions, which can be derived using equation (1.30) within the 3+1 framework, are given by [93]

$$\begin{aligned} P_{\nu_e \rightarrow \nu_e} &= 1 - 4 \sin^2 2\theta_{ee} \sin^2(1.27\Delta m^2 L/E) \\ P_{\nu_\mu \rightarrow \nu_\mu} &= 1 - 4 \sin^2 2\theta_{\mu\mu} \sin^2(1.27\Delta m^2 L/E) \\ P_{\nu_\mu \rightarrow \nu_e} &= 4 \sin^2 2\theta_{\mu e} \sin^2(1.27\Delta m^2 L/E), \end{aligned} \quad (1.36)$$

where Δm^2 , L , and E are in units of eV^2 , km, and GeV, respectively, and

$$\begin{aligned}\sin^2 2\theta_{ee} &= 4(1 - |U_{e4}|^2)|U_{e4}|^2 \\ \sin^2 2\theta_{\mu\mu} &= 4(1 - |U_{\mu4}|^2)|U_{\mu4}|^2 \\ \sin^2 2\theta_{\mu e} &= 4|U_{\mu4}|^2|U_{e4}|^2.\end{aligned}\tag{1.37}$$

The first expression in equation (1.36) can potentially explain the deficit of $\bar{\nu}_e$ and ν_e events observed in the RAA and gallium anomaly, respectively. Though both anomalies stem qualitatively from the same phenomenon— $\bar{\nu}_e$ disappearance at short baseline—the gallium anomaly in general prefers a larger value of $\sin^2 2\theta_{ee}$ than the RAA. This is evident in figure 1-13, which shows the regions in $\sin^2 2\theta_{ee}$ – Δm^2 parameter space preferred by the RAA and gallium anomalies, as well as global constraints from other experiments. These constraints come from short-to-medium-baseline reactor experiments, including NEOS [94], RENO [95], and Daya Bay [96], as well as very-short-baseline reactor experiments, including STEREO [97], DANSS [98], and PROSPECT [99]. Each of these experiments searches for $\bar{\nu}_e$ disappearance in a reactor-flux-agnostic way: the former through comparisons of the reactor $\bar{\nu}_e$ spectra measured by different detectors [100], and the latter through the use of modular or movable detectors capable of comparing $\bar{\nu}_e$ interaction rates across different baselines. The KATRIN experiment, which is sensitive to the neutrino mass via an extremely precise measurement of the tritium beta spectrum endpoint, also places strong constraints on $\sin^2 2\theta_{ee}$ in the $\Delta m^2 \gtrsim 10 \text{ eV}^2$ region [101].

The second expression in equation (1.36) can potentially explain the SBL anomalies. This is because both LSND and MiniBooNE operated at accelerator neutrino sources for which the neutrino flux was generated mainly by charged pion decay [5,14]; thus, due to helicity suppression, the flavor composition was dominated muon-flavored (anti)neutrinos. This means that even a small value of $\sin^2 2\theta_{\mu e}$ could generate an observable level of $\bar{\nu}_e$ appearance on top of the SM $\bar{\nu}_e$ flux prediction. Figure 1-14 shows the allowed regions in $\sin^2 2\theta_{\mu e}$ – Δm^2 parameter space from LSND and MiniBooNE [10]. Strikingly, both anomalies generally prefer the same region of parameter




Figure 1-13: Preferred regions in $\sin^2 2\theta_{ee}-\Delta m^2$ parameter space to explain the RAA [8] (green contour) and gallium anomaly [9] (blue regions). The total excluded region from other experiments (grey region) is also shown. Figure from Ref. [9].

space. However, as the MiniBooNE excess tends to peak more sharply at lower energies, the 3+1 fit prefers lower values of Δm^2 compared to the LSND result.

It is important to note that the fits performed in figure 1-14 account only for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations, ignoring any potential $\bar{\nu}_e$ or $\bar{\nu}_\mu$ disappearance in the SM background prediction. This is a reasonable approximation, however, the inclusion of the latter effects does indeed impact the MiniBooNE allowed regions. This effect was only accounted for recently in Ref. [102], which is presented in section 5.3.1 of this thesis.

While there are indications of short baseline $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ appearance and $\bar{\nu}_e$ disappearance in the global anomaly picture, direct observation of $\bar{\nu}_\mu$ disappearance via the third expression in equation (1.36) remains elusive. Long baseline experiments such as MINOS/MINOS+ [103, 104] and CCFR84 [105] have searched for muon neutrino disappearance from an accelerator neutrino source. Additionally, the IceCube experiment has searched for a sterile-induced matter resonance impacting muon neutrinos as they transit through the earth [106]. So far, no definitive evidence for $\bar{\nu}_\mu$ disappearance has been found (up to a $\sim 2\sigma$ preference in the IceCube results [106]).

The lack of $\bar{\nu}_\mu$ disappearance introduces significant tension when one tries to fit

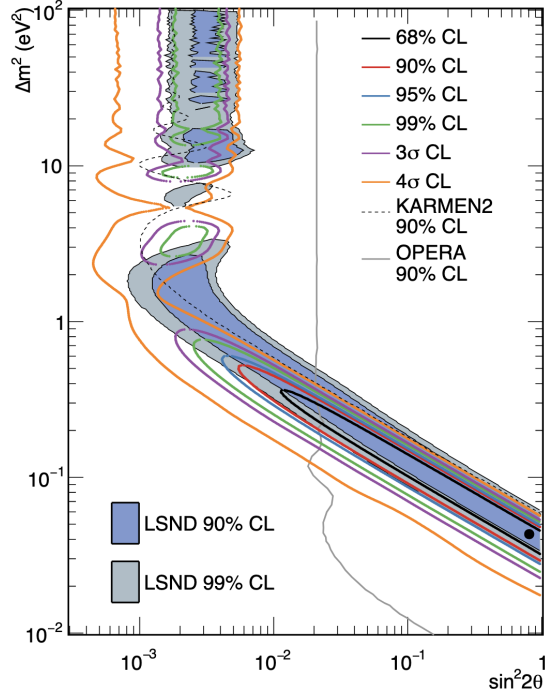




Figure 1-14: Preferred regions in $\sin^2 2\theta_{\mu e}-\Delta m^2$ parameter space to explain the LSND anomaly [5] (filled contours) and MiniBooNE anomaly [10] (open contours). Figure from Ref. [10].

global neutrino data within a consistent 3+1 model. This conclusion has been reached by multiple 3+1 global fitting efforts [11, 12, 93]; figure 1-15 shows a representation of the tension between appearance and disappearance experiments observed in global fits. This tension persists even with the inclusion of the recent BEST result, which prefers larger values of $|U_{e4}|^2$ (thus allowing lower values of $|U_{\mu 4}|^2$ to fit the $\bar{\nu}_e$ appearance anomalies) [12]. Thus, the 3+1 model, while still an important benchmark BSM scenario, has become disfavored as a solution to all observed anomalies in the neutrino sector. The state of the sterile neutrino explanation of the SBL anomalies is discussed in more detail throughout this thesis.

In recent years, neutrino physicists have begun to turn toward alternative explanations of the anomalies, often involving dark sector particles with additional interactions. Chapter 6 of this thesis covers one such explanation of the MiniBooNE anomaly, involving heavy right-handed neutrinos with a transition magnetic moment coupling to the active neutrinos.



(a) From Ref. [11]



(b) From Ref. [12]

Figure 1-15: Graphical representation of the tension observed in 3+1 global fits between different subsets of the experimental landscape. Figure 1-15a shows the tension between ν_e appearance experiments and ν_e/ν_μ disappearance experiments observed in Ref. [11]. Figure 1-15b shows the tension between allowed regions from ν_e appearance (lower right), ν_e disappearance (upper right), and ν_μ disappearance (upper left) experiments observed in Ref. [12], which includes the latest results from the BEST experiment.

Chapter 2

The MiniBooNE Experiment

This chapter is intended to give the reader an overview of the Mini Booster Neutrino Experiment (MiniBooNE), specifically concerning the excess of electron-like events observed by MiniBooNE in data taken between 2002–2019 at Fermilab’s Booster Neutrino Beam (BNB). The MiniBooNE excess is at the center of this thesis; the research presented here covers both experimental follow-up and theoretical interpretations of this anomaly. Thus, the remaining chapters require a thorough discussion of the MiniBooNE experiment and its most famous result.

2.1 Overview of MiniBooNE

MiniBooNE was originally designed as an experimental follow-up to the LSND excess of $\bar{\nu}_e$ events observed at the Los Alamos Meson Physics Facility (LAMPF) [5]. As described in section 1.4, the LAMPF flux comprised mostly of $\bar{\nu}_\mu$, which dominated over the $\bar{\nu}_e$ flux by three orders of magnitude [107]. Because of this, LSND was able to perform a low-background search for the IBD interaction $\bar{\nu}_e p \rightarrow e^+ n$. An excess of IBD events was observed above the intrinsic $\bar{\nu}_e$ flux prediction from the beam dump source—this is known as the “LSND anomaly” [5].

The LSND anomaly has traditionally been interpreted as evidence for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations at $\Delta m^2 \approx 1 \text{ eV}^2$. The LSND detector sat relatively close to the LAMPF nuclear target; the characteristic length-to-energy ratio in the experiment was $L/E \sim$

30 m/30 MeV. According to equation (1.32), in order to be sensitive to the oscillation-based interpretation of the LSND anomaly, one must maintain the same ratio L/E . This was the design strategy of the MiniBooNE experiment, which observed the interactions of neutrinos from the BNB with characteristic energy $E_\nu \sim 500$ MeV, at a baseline of $L \sim 500$ m from the BNB beryllium target. The BNB produced mostly ν_μ from pion decay-in-flight; thus, MiniBooNE searched for $\nu_\mu \rightarrow \nu_e$ oscillations in the BNB at $\Delta m^2 \approx 1$ eV².

2.1.1 The Booster Neutrino Beam

The BNB, which is still operational, follows the typical design of a neutrino beamline [107]. Protons are accelerated in a synchrotron up to a momentum of 8.89 GeV, at which point they are kicked out of the synchrotron and interact within the Be target of the BNB, producing a cascade of secondary particles [14]. The charged mesons in this cascade are then focused using a toroidal magnetic field from an aluminum horn. By switching the direction of the current in the horn (and thus the direction of the magnetic field), one can choose whether to focus positively-charged mesons and de-focus negatively-charged mesons (“neutrino mode”), or vice-versa (“antineutrino mode”). After focusing, charged mesons pass through a concrete collimator and enter a 50-meter-long air-filled region where they decay to neutrinos. The neutrinos travel through another 474 meters of bedrock before reaching the MiniBooNE detector. A schematic depiction of this process is shown in figure 2-1.

The MiniBooNE flux is described in detail in Ref. [14]; we summarize the most important details here. In neutrino (antineutrino) mode, the flux is dominated by ν_μ ($\bar{\nu}_\mu$) from π^+ (π^-) decay. Wrong-sign $\bar{\nu}_\mu$ (ν_μ), coming mostly from the decay of oppositely-charged pions, contribute at the 5% (15%) level. Two and three-body kaon decays also contribute to the ν_μ and $\bar{\nu}_\mu$ flux at the few-percent level. The BNB also produces electron (anti)neutrinos, which represent $< 1\%$ of the total flux in both neutrino and antineutrino mode. These come from two main sources: the decay of the secondary muon produced in the original charged pion decay, which is dominant for $E_\nu \lesssim 1$ GeV, and two-body kaon decay, which is dominant for $E_\nu \gtrsim 1$ GeV. The

BNB flux breakdown in neutrino and antineutrino mode are shown in figure 2-2.

The π^\pm production rate from p-Be interactions has been measured by the HARP [108] and BNL E910 [109] experiments. HARP took data at the BNB proton incident momentum (8.89 GeV/c) with a replica of the BNB beryllium target, while E910 took data at varying incident proton momenta above and below the nominal BNB value. These data were used to constrain a Sanford-Wang parameterization of the π^\pm differential production cross section in the BNB [110]. The charged and neutral kaon production rates in p-Be were constrained by measurements from other experiments at proton momenta around 8.89 GeV/c; the Feynman scaling hypothesis was used to relate these measurements to the BNB proton momentum [14].




Figure 2-1: A schematic depiction of the BNB at Fermilab, including the downstream MiniBooNE detector. Figure from Ref. [13].




Figure 2-2: Breakdown of the neutrino flux at the BNB in neutrino (left) and antineutrino (right) mode. Figure from Ref. [14]

2.1.2 The MiniBooNE Detector

The MiniBooNE detector is an 818-ton, 6.1-meter-radius spherical volume filled with mineral oil (approximately CH_2) [15]. It was designed to measure the Cherenkov

light produced from charged particles created in the charged-current interactions of BNB neutrinos within the detector volume. To do this, the inner surface of the sphere was instrumented with 1280 photo-multiplier tubes (PMTs), corresponding to a photocathode coverage of 11.3%. An additional 240 PMTs were used to instrument the surrounding veto region, which rejected cosmic muons and neutrino interactions outside the detector volume with an efficiency of $\sim 99.99\%$ [15]. Mineral oil was chosen as the detector medium due to its high index of refraction ($n=1.47$), leading to more Cherenkov light production by electrons traversing the detector volume. The exact mineral oil mixture, Marcol 7, was selected by optimizing the behavior of photons with wavelengths between 320 nm and 600 nm (e.g., requiring an extinction length greater than 20 m) [15]. The detector was situated in a cylindrical vault just below ground level, under ~ 3 m of dirt. A schematic depiction of the MiniBooNE detector is shown in figure 2-3.

The reconstruction of the final state from a neutrino interaction in MiniBooNE relied on the detection of Cherenkov light. Specifically, the collaboration developed reconstruction algorithms that turned the spatiotemporal distribution of photon hits on the PMTs into kinematic information on each observable final state particle [77]. These algorithms used maximum likelihood estimators to estimate the starting location, direction, and energy of final state particles using the observed photon hits in each PMT, relying on the known transport properties of Cherenkov photons within the detector medium. Cherenkov photons are emitted when a charged particle travels faster than the speed of light in a medium, at an angle of $\cos\theta_C = 1/n\beta$ with respect to the charged particle track. This results in a characteristic ring-like pattern on the detector wall. Such Cherenkov rings formed the basis of the MiniBooNE reconstruction algorithm.

There were two main classes for observable final state particles in MiniBooNE: muon-like (μ -like) and electron-like (e -like). Each elicits a different Cherenkov ring pattern [77]. At MiniBooNE energies, muons are typically minimum-ionizing particles and thus would appear as a uniform ring in the PMT array. The ring would be filled in if the muon exits the detector volume before going below the Cherenkov threshold, and

would be open otherwise. Electrons, on the other hand, undergo radiative processes as they travel, emitting photons via the Bremsstrahlung process, which then undergo pair-production to e^+e^- , which then emit more photons, and so on. This process is typically called an “electromagnetic (EM) shower”. The multiple constituent electrons and positrons in this EM shower would result in a distorted Cherenkov ring in the PMT array. Importantly, high energy photons also produced these distorted rings after undergoing an initial pair-production interaction; thus, electrons and photons were essentially indistinguishable in MiniBooNE and were both classified as e -like. Another relevant final-state particle in MiniBooNE was the neutral pion, which could be identified via two separate distorted Cherenkov rings via the $\pi^0 \rightarrow \gamma\gamma$ decay. It is also important to note that π^0 events could be misclassified as e -like if one of the photons was not reconstructed, which could happen if one of the photons escaped the detector without pair producing or had energy below the detection threshold. A schematic diagram of the detector signature of muons, electrons, and neutral pions in MiniBooNE is shown in figure 2-4a.

A separate likelihood was calculated for three different final state particle hypotheses: electron, muon, and neutral pion [77]. Ratios of these likelihoods were used to distinguish one particle from another in MiniBooNE. As an example, we show the separation of electron and muon events as characterized by the log-likelihood-ratio as a function of reconstructed neutrino energy in figure 2-4b. This ratio was the main selection tool in selecting e -like events for MiniBooNE’s flagship search for $\nu_\mu \rightarrow \nu_e$ oscillations in the BNB.

2.2 The MiniBooNE Low Energy Electron-Like Excess

As stated above, MiniBooNE was designed to test the LSND excess of $\bar{\nu}_e$ events discussed in section 1.4. To do this, MiniBooNE isolated a sample of e -like events using the likelihood ratios described in the previous section [76]. This sample was




Figure 2-3: The MiniBooNE detector situated in the cylindrical detector hall (left) and an image of the interior of the MiniBooNE detector (right), showing the PMTs in both the signal and veto regions. Figure from Ref. [15].

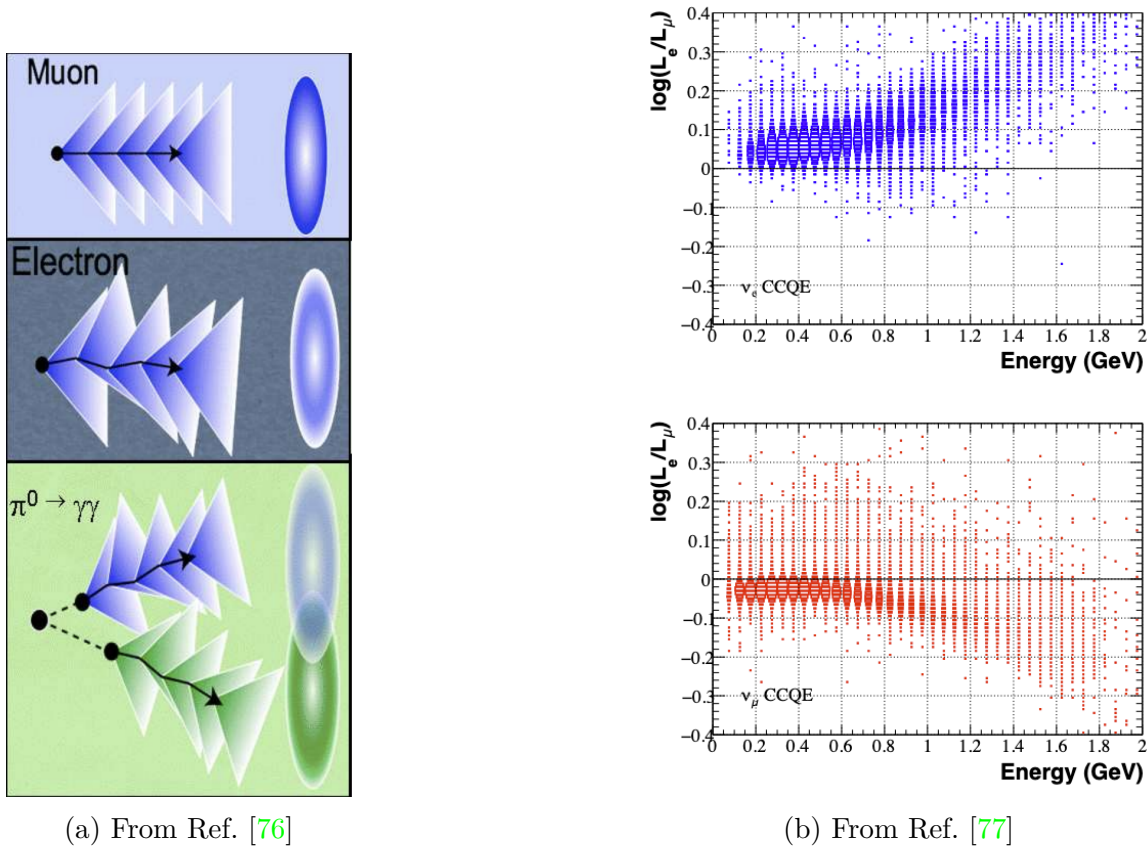


Figure 2-4: Visual representations of particle identification in MiniBooNE. Figure 2-4a shows a schematic representation of the detector signature from the three main particle classes in MiniBooNE: muons, electrons, and neutral pions. Figure 2-4b shows the MiniBooNE log-likelihood-ratio between the e -like and μ -like hypothesis as a function of reconstructed neutrino energy, considering both simulated ν_e CCQE (top) and ν_μ CCQE (bottom) interactions.

optimized to select ν_e CCQE interactions within the detector while rejecting ν_μ interaction backgrounds, thus maximizing sensitivity to potential $\nu_\mu \rightarrow \nu_e$ oscillations within the BNB. MiniBooNE’s flagship e -like analysis, which has remained stable over the lifetime of the experiment, achieved a ν_e CCQE efficiency of $\sim 20\%$ while rejecting $\sim 99.9\%$ of ν_μ backgrounds [10]. The full MiniBooNE dataset consists of 18.75×10^{20} (11.27×10^{20}) protons-on-target (POT) in neutrino (antineutrino) mode collected over 17 years of operation. In this dataset, the e -like analysis observes 2870 (478) data events in neutrino (antineutrino) mode, compared to an SM prediction of 2309.4 (400.6) events [10]. Combining neutrino and antineutrino mode data, MiniBooNE observes an excess of 638.0 ± 52.1 (stat.) ± 122.2 (sys.) e -like events, corresponding to a Gaussian significance of 4.8σ [10].

Figure 2-5 shows the reconstructed neutrino energy distribution of the MiniBooNE e -like excess in both neutrino and antineutrino mode. The stacked histogram corresponds to the SM prediction from the NUANCE event generator [44], while the data points correspond to the observed number of e -like events in each bin. The error bars on the stacked histogram correspond to the systematic uncertainty on the SM prediction, calculated within a covariance matrix formalism. The dominant sources of systematic uncertainty include neutrino cross section modeling (derived largely using MiniBooNE’s own cross section measurements [111–114]), nuclear effects, detector response and optical modeling, and BNB flux estimation [10, 79]. The presented error in each bin of the ν_e and $\bar{\nu}_e$ sample incorporates a constraint from MiniBooNE’s dedicated ν_μ and $\bar{\nu}_\mu$ CCQE samples. The dashed line corresponds to the best fit of the MiniBooNE excess to the $3+1$ sterile neutrino model described in section 1.4. As one can see, the excess in data events is strongest in the lowest energy bins; for this reason, this anomaly is often referred to as the MiniBooNE low-energy excess (LEE).

As MiniBooNE used a Cherenkov detector, it was not sensitive to the final state hadronic activity in a neutrino interaction. Thus, kinematic reconstruction of the original neutrino relied entirely on the final state lepton. Under the assumption that the neutrino underwent a CCQE interaction off of a nucleon at rest within the nucleus,

the original neutrino energy is given by [111]

$$E_\nu^{\text{QE}} = \frac{2(M'_n)E_\ell - ((M'_n)^2 + m_\ell^2 - M_p)}{2[(M'_n) - E_\ell + \sqrt{E_\ell^2 - m_\ell^2} \cos \theta_\ell]}, \quad (2.1)$$

where E_ℓ is the total lepton energy, $\cos \theta_\ell$ is the lepton scattering angle, and M_n , M_p , and m_ℓ are the neutron, proton, and lepton mass, respectively. The adjusted neutron energy is defined as $M'_n \equiv M_n - E_B$, where E_B is the nuclear binding energy of the initial state neutron. An analogous relation exists for antineutrino energy reconstruction in a CCQE interaction [115]. This is the reconstructed energy definition used to generate the histograms in figure 2-5.

Figure 2-6 shows the visible energy and $\cos \theta$ distributions of the final state lepton in MiniBooNE's e -like neutrino mode sample. The visible energy distribution shows the strongest discrepancy for softer lepton kinetic energies, as expected for a low-energy excess. For the angular distribution, it is worth noting that while there is an excess across the entire range, the largest deviation above the SM prediction comes from the $\cos \theta \in [0.9, 1.0]$ bin. The angular distribution of the MiniBooNE LEE is an important piece of information for potential solutions to the anomaly—as we will discuss throughout this thesis, BSM physics models often cannot explain the energy and angular distributions of the MiniBooNE LEE simultaneously.

The green contributions to the stacked histograms of figure 2-5 represent the interactions of intrinsic ν_e or $\bar{\nu}_e$ in the BNB. At low energies, these events come mostly from the decay of the secondary muon in the $\pi^+ \rightarrow \mu^+$ or $\pi^- \rightarrow \mu^-$ decay chain, while ν_e and $\bar{\nu}_e$ from kaon decays start to contribute more at higher energies [14].

The red and brown contributions to the stacked histograms represent misidentified photon backgrounds that are reconstructed as a single distorted Cherenkov ring. The largest photon background comes from misidentified π^0 created via ν_μ and $\bar{\nu}_\mu$ neutral-current (NC) resonant scattering, in which the initial state nucleon is excited to a Δ resonance before decaying to nucleon-pion pair. The π^0 decay promptly to a pair of photons, which should nominally appear as a pair of distorted Cherenkov rings as in figure 2-4a. However, if one of the photons exits the detector before converting to


an e^+e^- pair, or if the original pion energy is distributed asymmetrically such that the visible energy of one of the photons sits below the reconstruction threshold of 140 MeV [76], the π^0 decay will be misidentified as an e -like event. An enhancement of the NC π^0 background in figure 2-5 could in principle explain the observed excess. However, MiniBooNE constrained the rate of this NC π^0 background *in situ* via a measurement of the two-gamma invariant mass peak in well-reconstructed NC π^0 events [10]. Additionally, the radial distribution of the excess peaks toward the center of the detector, while misidentified NC π^0 backgrounds happen more often toward the edge of the detector, where it is more likely for a photon to escape before pair-producing.

The next-largest photon background comes from rare $\Delta \rightarrow N\gamma$ decays in ν_μ and $\bar{\nu}_\mu$ NC resonant scattering interactions. As this process has never been observed directly, it was not possible for MiniBooNE to constrain the $\Delta \rightarrow N\gamma$ event rate *in situ*. It was instead constrained indirectly by the NC π^0 two-gamma invariant mass distribution [10]. A factor 3.18 enhancement in $\Delta \rightarrow N\gamma$ events could explain the MiniBooNE LEE; however, this hypothesis has since been disfavored by recent results from the MicroBooNE experiment [80]. The MicroBooNE experiment will be covered in more detail in chapters 3 to 5.


MiniBooNE has also studied neutrino interactions outside the detector volume which result in a single photon entering the detector (the “dirt” backgrounds in figure 2-5). The timing distribution of the MiniBooNE e -like dataset suggests that the excess comes primarily in time with the beam, while dirt background events are often delayed by ~ 10 ns [10]. This result disfavors an enhancement of external neutrino interactions as an explanation of the MiniBooNE excess.

Therefore, the 4.8σ MiniBooNE excess remains unexplained. Resolution of the MiniBooNE LEE is one of the major goals of the neutrino community [83].

The MiniBooNE LEE is most commonly interpreted within the context of the 3+1 model introduced in section 1.4. This is primarily because the MiniBooNE excess has historically been considered alongside the LSND excess, as both results can be explained by short-baseline $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ appearance. Strikingly,

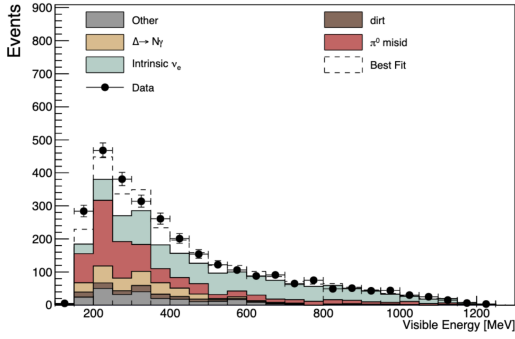


(a) ν_e sample, from Ref. [10]




(b) $\bar{\nu}_e$ sample, from Ref. [79]

Figure 2-5: The E_ν^{QE} distribution of the MiniBooNE e -like excess in the total neutrino mode (figure 2-5a) and antineutrino mode (figure 2-5b) datasets. The observation and SM prediction in each bin are shown by the data points and colored histograms, respectively.



(a) Lepton E_{vis} distribution



(b) Lepton $\cos \theta$ distribution

Figure 2-6: The lepton visible energy (figure 2-6a) and $\cos \theta$ (figure 2-6b) distributions of the MiniBooNE e -like excess in the total neutrino mode dataset. The observation and SM prediction in each bin are shown by the data points and colored histograms, respectively. Figures from Ref. [10].

the MiniBooNE and LSND anomalies both prefer similar regions in sterile neutrino parameter space, as shown in figure 1-14. This is further supported by figure 2-7, which shows the rising nature of the MiniBooNE and LSND excesses as a function of the ratio L/E , behavior that is consistent with a sterile neutrino explanation.

There are, however, complications regarding a sterile neutrino explanation of the MiniBooNE excess. The 3+1 model has difficulty reproducing the lowest energy and lowest scattering angle region of the excess. Figure 2-6 shows that the best fit 3+1 prediction, indicated by the dotted line, still falls below the observed data in

the lowest lepton E_{vis} bin and highest lepton $\cos\theta$ bin. Additionally, as discussed in section 1.4, there is significant tension between the MiniBooNE and LSND observation of $\nu_e/\bar{\nu}_e$ appearance and experiments searching for $\nu_e/\bar{\nu}_e$ and $\nu_\mu/\bar{\nu}_\mu$ disappearance. Finally, the follow-up MicroBooNE experiment has not observed an excess of ν_e events consistent with the expectation from the MiniBooNE LEE [116]. The MicroBooNE ν_e analysis is one of the main results of this thesis and will be explored in more detail in chapters 4 and 5. While the non-observation of a MiniBooNE-like excess of ν_e events in the BNB does set constraints on $3 + 1$ parameter space, it does not fully exclude the MiniBooNE allowed regions [102]. This point will be discussed further in chapter 5.

These complications with the eV-scale sterile neutrino interpretation of the MiniBooNE LEE have prompted the community to explore alternative explanations. Many of these are relatively simple extensions beyond $3 + 1$, such as $3 + N$ models involving N additional sterile neutrino states [93], decaying sterile neutrino models [117–119], and sterile neutrinos with altered dispersion relations from large extra dimensions [120–122]. Other explanations for the MiniBooNE LEE introduce a number of new particle species prescribed with new interactions that create an additional source of photons or e^+e^- pairs in MiniBooNE. Such models include heavy neutral leptons (HNLs) which decay to photons via a transition magnetic moment [27, 31, 123–135] and models with heavy neutrinos coupled to a “dark sector” involving, for example, new vector or scalar mediators [136–146]. Chapter 6 of this thesis explores one such explanation of MiniBooNE involving an HNL with a transition magnetic moment coupling to active neutrinos, which we hereby refer to as a “neutrissimo”. Neutrissimo decays in MiniBooNE provide an additional source of single photons which could explain the e -like excess [27, 31].

Thus, there are many potential explanations for the MiniBooNE anomaly. Distinguishing between these explanations requires careful consideration of the kinematic distributions of the MiniBooNE excess [31, 147]. Further, these models are often subject to constraints from existing accelerator neutrino experiments, such as MINERvA and NA62 [139, 146]. A complete evaluation of constraints from existing data




Figure 2-7: The MiniBooNE and LSND excesses as a function of the ratio L/E . The MiniBooNE data is separated into neutrino and antineutrino mode. Figure from Ref. [10].

is essential in determining the most viable models among the many proposed MiniBooNE explanations. The work presented in chapter 6 takes a step in this direction by calculating constraints from MINERvA on the neutrissimo model.

Chapter 3

The MicroBooNE Detector

In order to ascertain the nature of the MiniBooNE excess, one needs a detector capable of providing more detailed event-by-event information than MiniBooNE’s Cherenkov detector. This is the concept behind the MicroBooNE experiment. The MicroBooNE detector is a large-scale liquid argon time projection chamber (LArTPC) with the ability to record high-resolution images of neutrino interactions. MicroBooNE recently released its first results investigating the nature of the MiniBooNE excess [80, 116], which will be presented in chapters 4 and 5. This chapter introduces the detector that made this measurement possible.

3.1 Liquid Argon Time Projection Chamber

MicroBooNE used an 85-metric-ton fiducial volume LArTPC detector to observe the interactions of neutrinos in the BNB [16, 148]. This makes MicroBooNE the first $\mathcal{O}(100\text{ t})$ LArTPC operated in the United States. The idea for a LAr-based total absorption detector originated in the 1970s [149]. The introduction of the LArTPC detector concept came from Carlo Rubbia in 1977 [150], extending earlier work from David Nygren [151] and Georges Charpak [152]. The first operational large-scale LArTPC was the 500-metric-ton active volume ICARUS T600 detector [153], which came online in 2010. ICARUS observed cosmic ray and neutrino interactions at the Gran Sasso underground National Laboratory [154] and even set constraints on

$\nu_\mu \rightarrow \nu_e$ interpretations of the LSND and MiniBooNE anomalies using the CERN to Gran Sasso neutrino beam [155]. On a smaller scale, the ArgoNeuT experiment operated a 0.25-metric-ton LArTPC at Fermilab’s Neutrino Main Injector beamline from 2009-2010, where it performed the first measurements of neutrino-argon cross sections [156].

The MicroBooNE detector is situated 70 m downstream of the MiniBooNE detector along the BNB and operated from 2015 to 2021, observing a total of approximately 1.5×10^{21} POT [148]. MicroBooNE LArTPC data come in the form of high-resolution three-dimensional images of the ionization energy deposited by final state charged particles in neutrino interactions. The information contained in these images allows for the event-by-event separation of photons and electrons—an essential capability for determining the source of the MiniBooNE excess. MicroBooNE can also reconstruct hadronic activity in the final state of the neutrino interaction, which helps further distinguish between the possible sources of the MiniBooNE excess.

We begin with a brief overview of the MicroBooNE reconstruction procedure. Charged-current neutrino interactions in the LAr volume produce charged particles in the final state, which ionize argon atoms as they traverse the detector. Thus, each charged particle leaves behind a trail of ionized electrons which, in theory, can drift freely through the noble element detector medium without being captured. This drift is controlled via an external electric field with strength $|E| \sim 273$ V/cm, accelerating the ionized electrons to a final velocity $v \sim 0.11$ cm/ μ s toward three anode wire planes [148]. MicroBooNE employs a right-handed coordinate system, in which BNB neutrinos travel along the \hat{z} direction, ionization electrons drift along the $-\hat{x}$ direction, and \hat{y} represents the vertical direction [16]. The anode planes consist of two induction planes and one collection plane, each containing a series of wires spaced 3 mm apart and oriented at $\pm 60^\circ$ and 0° with respect to the \hat{y} direction for the induction and collection planes, respectively. Each plane is biased such that ionization electrons drift past the induction plane wires, generating a signal via induction, and terminate on the collection plane wires, generating a signal via direct charge collection. The signals on the anode wire planes allow for two-dimensional reconstruction of the charged particle

trajectory in the $\hat{y} - \hat{z}$ plane, transverse to the drift direction. The \hat{x} dimension of the charged particle trajectory can be reconstructed using the arrival time of signals on the anode wires in conjunction with the known drift time of the ionization electrons. In order for this technique to work, one must know the initial time at which the charged particle entered the detector. This can be established using either an external beam trigger or an internal trigger from the light collection system, which operates on much shorter time scales, $\mathcal{O}(\text{ns})$, compared to characteristic electron drift times of $\mathcal{O}(\text{ms})$. A schematic of this process is shown in figure 3-1.

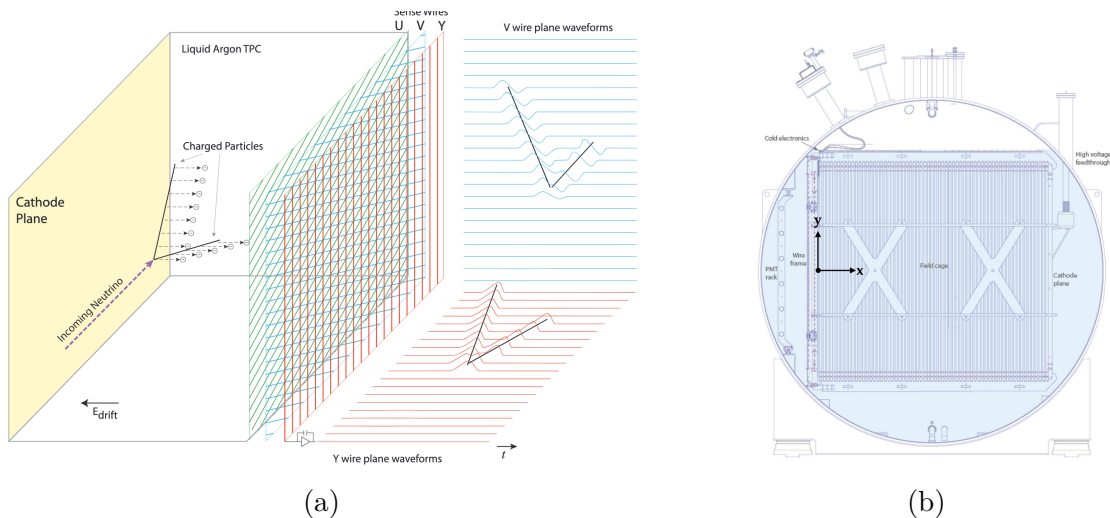


Figure 3-1: Schematic depictions of the MicroBooNE LArTPC. Figure 3-1a shows the detection process for charged particles from a neutrino interaction in a MicroBooNE-like LArTPC. Figure 3-1b shows a cross-sectional view of the MicroBooNE detector along the $-\hat{z}$ direction. Figures from Ref. [16].

3.1.1 Cryogenics

The MicroBooNE detector is relatively large—the LArTPC volume spans 2.6 m, 2.3 m, and 10.4 m in the \hat{x} , \hat{y} , and \hat{z} direction, respectively [16]. Thus, ionization electrons must drift through $\mathcal{O}(m)$ of LAr before reaching the anode wire planes. Reconstruction of these ionization electrons requires careful control of the drift process. This is the main objective of the MicroBooNE cryogenic system.

The LArTPC is housed within a larger cylindrical cryostat, which itself is supported by an argon purification system and nitrogen refrigeration system [16]. The

purification system consists of two recirculation pumps and two filter skids that remove electronegative impurities from the LAr, mainly oxygen (O_2) and water (H_2O). These impurities must be kept below the 100 parts-per-trillion O_2 -equivalent level in order to maintain electron drift lengths of at least 2.5 m [16,157]. Additionally, the nitrogen contamination must be kept below 2 parts-per-million in order to maintain an argon scintillation light attenuation length greater than the size of the detector [158]. Nitrogen cannot be appreciably removed from the argon via the purification system; rather, the initial nitrogen contamination is fixed by the quality of the delivered argon, and additional contamination must be controlled by minimizing the atmosphere leakage rate into the cryostat.

The nitrogen refrigeration system is designed to combat the heat load on the LAr from the environment and electrical power systems, maintaining thermal homogeneity throughout the active volume. It consists of two condensers, each designed to handle a heat load of approximately 9.5 kW [16]. The temperature of the LAr volume must be stable to ± 0.1 K in order to keep the \hat{x} direction resolution of charged particle tracks below 0.1% [16].

3.1.2 LArTPC Drift System

The drift system inside the LArTPC volume consists of three major subsystems: the cathode plane, the field cage, and the three anode wire planes. The purpose of the drift system is to maintain a uniform electric field throughout the active volume such that ionization electrons are transported to the anode plane at a stable drift velocity.

The cathode consists of nine stainless steel sheets connected to a supporting frame to form a single plane. Laser tracker measurements indicate that a majority of the cathode plan is flat to within ± 3 mm [16]. The cathode plane is kept at a negative potential of approximately -70 kV via a high voltage feedthrough on the cryostat. The field cage maintains a uniform electric field between the cathode plane and anode planes. It consists of 64 stainless steel tubes wrapped around the LArTPC active volume. A resistor divider chain connects each tube to its neighbor, sequentially stepping the voltage from -70 kV to ground in 1.1 kV increments. The chain pro-

vides a resistance of $250\text{ M}\Omega$ between adjacent tubes such that the current flow is approximately $4.4\text{ }\mu\text{A}$, much larger than the $\mathcal{O}(\text{nA})$ current from signals on anode plane wires [16]. Figure 3-2 shows the MicroBooNE cathode and field cage, as well as a simulated map of the electric field within the LArTPC active volume.


Perhaps the most critical components of the MicroBooNE detector are the three anode wire planes. The U and V induction planes contain 2400 wires each, while the Y collection plane contains 3456 wires. As mentioned above, the U and V plane wires are oriented at $\pm 60^\circ$ with respect to the vertical, while the Y plane is oriented vertically. The U, V, and Y planes are biased at -200 V , 0 V , and $+440\text{ V}$, respectively, to ensure termination of ionization electrons on the Y collection plane. Each wire is $150\text{ }\mu\text{m}$ in diameter and is spaced 3 mm from its neighbors. The planes themselves are spaced 3 mm from one another. The wires are held in place by wire carrier boards, which house 16 wires each in the U and V planes and 32 wires in the Y plane. Each wire is terminated using a semi-automated wrapping procedure around a 3 mm diameter brass ferrule. On the wire carrier boards, each wire makes contact with a gold pin that connects to the electronic read-out system. The anode planes are held in place by a single stainless steel frame, which houses each wire carrier board via an array of precision alignment pins. Wires are tested to withstand three times the nominal load of 0.7 kg without breakage, both before and after placement onto the wire carrier board. Figure 3-3a shows an image of a single Y plane wire carrier board with 32 mounted wires. An image of the fully-assembled MicroBooNE LArTPC is shown in figure 3-3b, specifically highlighting the anode planes mounted on the stainless steel frame.

3.1.3 Light Collection System

Liquid argon is a prolific scintillation medium due to its low cost, high scintillation yield ($\mathcal{O}(10^4)$ photons per MeV of deposited energy), and transparency to its own scintillation light [158]. This last feature comes from the scintillation mechanism in LAr: when argon atoms are ionized, they combine with one another to form singlet and triplet excimer states. When these excimer states decay, they emit 128 nm




(a)



(b)

Figure 3-2: Figure 3-2a shows a close-up image of the cathode plane of the Micro-BooNE LArTPC. The stainless steel field cage tubes can also be seen surrounding the active volume. Figure 3-2b shows a cross-sectional map of the electric field at the edge of the active volume, considering a cathode plane voltage of -128 kV. The legend shows the field strength in units of V/m. Figures from Ref. [16].



(a)



(b)

Figure 3-3: Figure 3-3a shows a photograph of a single wire carrier board with 32 mounted wires. Figure 3-3b shows the fully-assembled MicroBooNE LArTPC, highlighting the anode plane mounted on the stainless steel frame. Figures from Ref. [16].

photons which pass unattenuated through the surrounding atomic argon [159]. The decay of the singlet (triplet) state happens on timescales of $\mathcal{O}(\text{ns})$ ($\mathcal{O}(\mu\text{s})$) [160, 161].

Thus, scintillation light emission happens on much shorter timescales than the $\mathcal{O}(\text{ms})$ drift time of the ionization electrons.

The light collection system in MicroBooNE is designed to detect the scintillation photons produced in a neutrino interaction. It consists of 32 8-inch Hamamatsu R5912-02mod cryogenic PMTs situated behind an acrylic plate coated with tetraphenyl butadiene (TPB) [16]. An image of one such PMT assembly is shown in figure 3-5a. TPB is a wavelength shifter that absorbs the 128 nm argon scintillation light and re-emits a photon in the visible range. The necessity of this procedure is shown in figure 3-4, which demonstrates that, unlike direct LAr scintillation light, TPB emission is well within the wavelength acceptance range of the PMTs.

When a photon hits the bi-alkali photocathode surface of an R5912-02mod PMT, an electron is released via the photoelectric effect [162]—this electron is often referred to as a “photoelectron” (p.e.). Each p.e. is focused toward a dynode chain—a series of electrodes designed to produce a number $n > 1$ electrons for each incident electron, resulting in an avalanche of electrons by the final anode which can be read out in the form of a current. The wavelength-dependent probability with which a given p.e. enters the dynode chain is known as the “quantum efficiency” of the PMT. At the nominal operating temperature of 87 K, the MicroBooNE PMTs have an average quantum efficiency of 15.3% [16]. Each PMT was tested in a liquid nitrogen (77 K) cryogenic environment, in which the gain and rate of thermal emission (“dark current”) of the PMT were measured as a function of the supplied high voltage (HV) across the dynode chain [163]. The HV for each PMT was set to produce a gain of 3×10^7 at 77 K [16], corresponding to a dark current of $\mathcal{O}(\text{kHz})$ [163]. The current output from each PMT passed through preamp/shaper boards before being digitized via an analog-to-digital converter (ADC) with a sampling rate of 64 MHz [16]. Thus, light is collected in time ticks with a length of 15.625 ns [148].

The light collection system was critical in detecting activity in the detector coincident with a beam spill, indicating the presence of a neutrino interaction. In order to record a given event, at least 5 photoelectrons must have been detected across all PMTs [148]. Additionally, a “common optical filter” was applied to reduce the

non-neutrino trigger rate—this filter required at least one string of six time ticks with greater than 20 photoelectrons detected during the 1.6 μs beam spill window, and no such strings of six time ticks in the 2 μs prior to the beam spill. The spatiotemporal distribution of light observed by the PMTs can also be used to augment the signal from the wire planes. For example, the PMT signal from a cosmic muon which stops in the detector and decays to a Michel electron is shown in figure 3-5b.

The scintillation light yield and detection efficiency in liquid argon are impacted by a number of phenomena. Impurities can dissociate the argon excimers before they have a chance to emit scintillation light. For example, O_2 molecules in the LAr volume can undergo a two-body collision with an argon excimer [164],



This interaction mainly decreases the decay probability of the longer-lived triplet state. Because of this, measurements of the delayed scintillation lifetime in LAr are sensitive to the concentration of impurities in the detector [16, 160, 164].

Scintillation light can also be absorbed by impurities within the detector. At concentrations of around 2 ppm, dissolved nitrogen will decrease the attenuation length of 128 nm scintillation light to around 30 m [158, 165]. TPB emanation from the painted acrylic plates can also lead to a bulk fluorescence effect within the liquid argon [166]. Additionally, Rayleigh scattering can deflect scintillation photons, diminishing the detector’s capability to translate photon detection into spatial information regarding the path of the original charged particle. The Rayleigh scattering length for 128 nm photons in liquid argon has been measured to be approximately 55 cm [167]. Given the size of the MicroBooNE detector, scintillation photons will undergo around five Rayleigh scattering interactions on average before reaching a PMT.

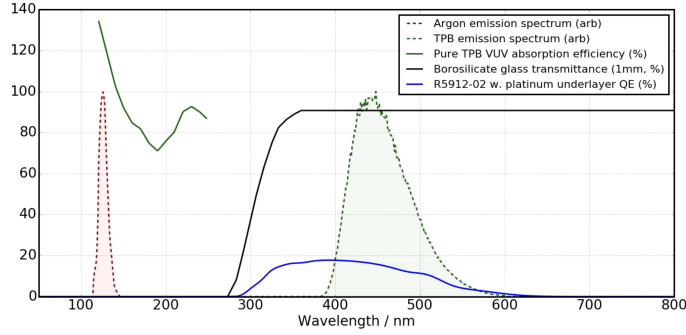

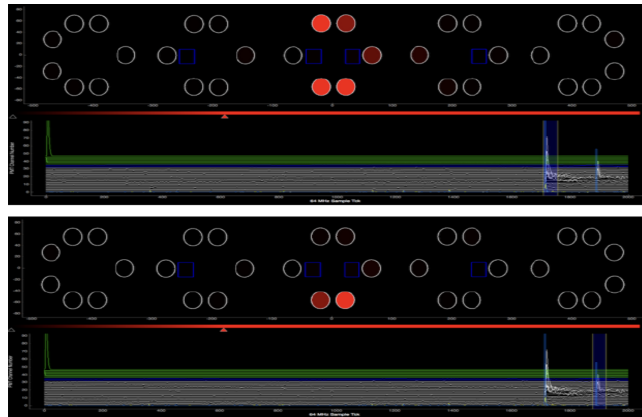


Figure 3-4: The LAr scintillation light spectrum, TPB ultra-violet absorption spectrum, TPB emission spectrum, PMT quantum efficiency, and PMT surface transmission efficiency as a function of photon wavelength. Figure from Ref. [17].



(a)



(b)

Figure 3-5: Figure 3-5a shows a photograph of a single PMT system used in the MicroBooNE detector. The acrylic window here has not yet been coated with TPB. Figure 3-5b shows the PMT signal from a stopped cosmic muon (top) that decays to a Michel (bottom). Figures from Ref. [16].

3.2 TPC Signal Processing

MicroBooNE has developed a custom cryogenic electronics system to amplify and digitize the analog signal on each wire of the LArTPC [16]. The wire signals first pass through front-end ASICs located on cryogenic motherboards attached directly to the wire carrier boards on the LArTPC. Each ASIC amplifies and shapes the signal from 16 wires—thus, 516 ASICs are needed to fully instrument the 8,256 readout channels of the LArTPC. Cryogenic-tested cables carry signals from the ASICs to intermediate amplifiers directly outside the cryostat, after which the signals travel

through 20 m cables to the readout system. This system consists of 130 readout modules, each containing 8 custom-designed 12-bit ADCs and a Front End Module (FEM). Each ADC digitizes the signal from 8 wires at a rate of 16 MHz. The dynamic range of the ADC is set to enable the detection both of a minimum-ionizing particle far away from the anode planes and of a highly-ionizing-particle close to the anode planes without saturation. The FEM uses a field-programmable gate array (FPGA) to store and compress the digitized signals. Once a trigger is received, the FPGA sends 4.8 ms of TPC data to the data acquisition (DAQ) system. The 4.8 ms window is digitized into 9600 time ticks on each wire, with each time tick corresponding to 500 ns of integrated charge [148].

3.2.1 Noise Filtering

The MicroBooNE LArTPC underwent an initial engineering run from October 2015 to July 2016. Various sources of noise in the TPC were identified using data from this run, all of which are described in detail in Ref. [18]. The frequency dependence of the noise was studied to inform the creation of an offline noise filter.

Some of this noise is inherent to the electronics system, coming from the cold, front-end ASIC or the two RC circuits in the intermediate amplifier and ADC board. The gain and peaking time of the front-end ASIC are configurable from among four pre-defined settings each, which determine the dynamic range and timing granularity of the ASIC. This choice also fixes the level of irreducible noise in the ASIC, comprised of multiple sources with different frequency dependence [18]. The timing response of the RC circuits can distort large amplitude, long duration signals such as a cosmic muon traveling parallel to a single wire. This effect can be corrected via a deconvolution process [18].

The engineering run also revealed a number of excess noise sources above the inherent electronics noise. The most prominent of these was low frequency ($\lesssim 30$ kHz) noise from the low voltage regulators on the front-end ASICs. Harmonic noise was also observed corresponding to the ripple frequency of the HV power supply for the field cage. Offline filters were developed to subtract the noise induced from both of

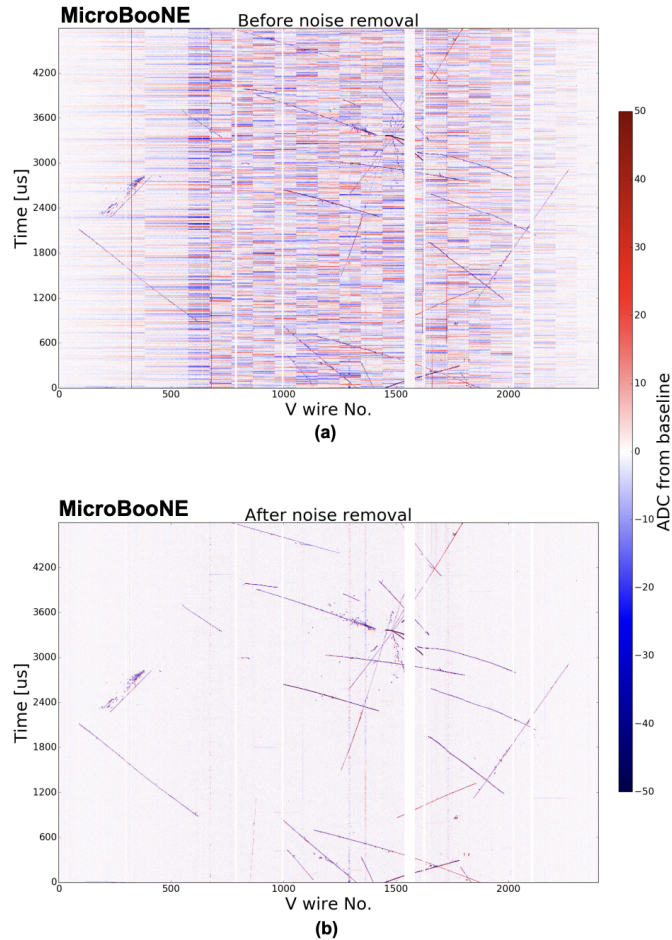


Figure 3-6: 2D displays of the signal from wires in one of the induction planes in a single data event, before and after the application of the offline noise filters. Figure from Ref. [18].

these sources. As shown in figure 3-6, these filters are able to remove most of the noise in the wire planes while having minimal impact on the true signal from ionization electrons. After the engineering run, hardware upgrades were made to the low voltage regulators and the HV power supply which greatly reduced the excess noise in the TPC readout signals. Even after these upgrades, the offline noise filters were still applied to optimize the TPC noise level.

3.2.2 Deconvolution

Even after noise filtering, the signal from the TPC is a complex combination of a number of factors: the ionization electron distribution, the induced current in each

wire due to passing ionization electrons, and the response of the electronics system. The signal must undergo deconvolution in order to obtain the quantity of physical interest—the number of ionization electrons passing each wire as a function of time. The signal processing algorithm designed to extract this information is described in detail in Ref. [19] and validated using MicroBooNE data in Ref. [168]. We give a brief overview of the algorithm here.

The measured signal $M(t)$ on a given wire is a convolution of the original signal $S(t')$ and the detector response function $R(t, t')$, which is intended to capture all the factors impacting signal formation discussed above. This can be expressed by the integral

$$M(t) = \int_{-\infty}^{\infty} R(t, t')S(t')dt'. \quad (3.2)$$

For a time-invariant detector response function $R(t, t') = R(t - t')$, one can express this relationship in the frequency domain via a Fourier transform. In a realistic setup, however, this treatment tends to amplify the effect of high-frequency components of the inherent electronics noise. Typically one introduces a filter function $F(\omega)$ to suppress high frequency response, such that the original signal can be determined from

$$S(\omega) = \frac{M(\omega)}{R(\omega)}F(\omega). \quad (3.3)$$

In this setup, one can consider $F(\omega)$ as a replacement for the detector response function.

The filter function used in MicroBooNE is inspired by the Wiener filter [169], which is designed to optimize the signal-to-noise ratio while minimizing variance and bias and the resulting $S(\omega)$. The exact Wiener filter is not applicable in the realistic MicroBooNE setup, as it does not conserve the total number of ionization electrons and can lead to non-local charge smearing due to suppression of low frequencies [19]. MicroBooNE instead constructs a Wiener filter $F(\omega)$ from simulation, and then fits the result to a functional form designed to preserve the normalization of the measured signal and remove the unwanted low-frequency suppression behavior.

The procedure discussed so far only accounts for the signal generated in a given

wire due to ionization electrons passing by that wire. In reality, electrons passing by neighboring wires can also induce a current in a given wire. This is addressed by using a two-dimensional (2D) deconvolution in both the time and wire dimensions rather than the one-dimensional deconvolution in the time dimension introduced above. An additional Weiner-inspired filter is constructed along the wire dimension when performing 2D deconvolution. After the 2D deconvolution, the algorithm identifies a region of interest (ROI) around each potential signal. These ROIs act as high-pass filters which suppress low-frequency noise in the deconvolved signal while preventing any non-local charge smearing.

The full deconvolution algorithm results in a robust reconstruction of the amount of charge passing by or deposited on each wire of the LArTPC. Figure 3-7 shows the U plane event display for a neutrino interaction candidate in MicroBooNE data after the application of the deconvolution algorithm. After the 2D deconvolution, one can clearly identify all of the charged particle tracks coming from the neutrino interaction vertex. The shape of the deconvolved signal for wires in each plane matches closely between data and simulation, as discussed in Ref. [168]. The signals produced by the deconvolution algorithm are ready to be used by MicroBooNE physics analyses, including the search for the MiniBooNE electron-like excess presented in chapters 4 and 5.

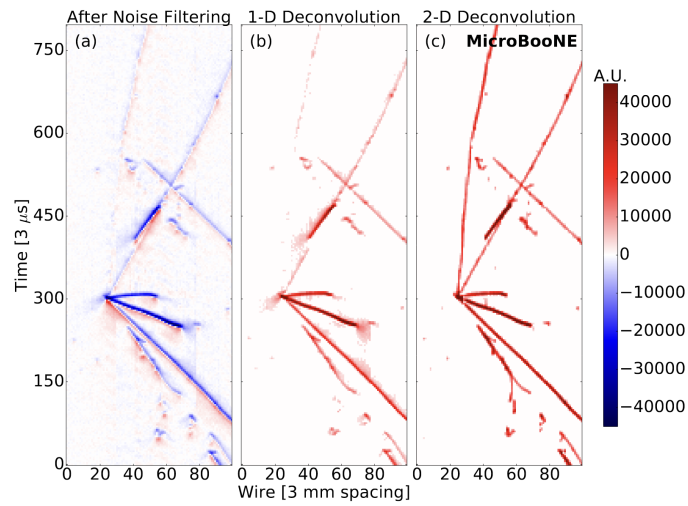


Figure 3-7: A U plane event display of a candidate neutrino interaction in the MicroBooNE data. The impact of the 1D and 2D deconvolution algorithms on the post-noise-filtering signal is shown. Figure from Ref. [19].

Chapter 4

The MicroBooNE Electron Neutrino Analysis: Overview and Selection

The MicroBooNE experiment was conceived to investigate the nature of the excess of low-energy electron-like events observed by the MiniBooNE experiment, described in detail in chapter 2. The MicroBooNE LArTPC, described in chapter 3, produces high-resolution event displays of neutrino interactions which enable the separation of events with electrons in the final state from those with photons. The collaboration released its first results searching for a MiniBooNE-like excess of electromagnetic events in October 2021. Four different analyses were performed, one focusing on photons and three focusing on electrons. The photon analysis searched for an excess of $\Delta \rightarrow N\gamma$ radiative decays [80]—the only SM photon background not constrained *in situ* by MiniBooNE. The three electron analyses searched for an excess of ν_e charged-current (CC) interactions, each with a different final state signal definition [116]. One of the electron analyses, hereafter referred to as the “inclusive analysis”, isolated all ν_e CC interactions using the WireCell reconstruction framework [170]. Another electron analysis, hereafter referred to as the “MiniBooNE-like analysis”, used the Pandora reconstruction framework to isolate all ν_e CC interactions without pions in the final state [171]. This chapter describes the “two-body CCQE analysis”, which used deep

learning techniques to isolate charged-current quasi-elastic (CCQE) ν_e interactions consistent with two-body scattering [148]. The results from this analysis are presented in chapter 5. The two-body CCQE analysis represents much of my early work as a graduate student.

Publications covered in this chapter for which I either held a leading role or made major contributions: [116, 148, 172]

4.1 Dataset and Simulation

The results covered in this section come from the first 6.88×10^{20} POT of BNB data taken by the MicroBooNE experiment. These data were separated into three run periods which spanned from 2016 to 2018: run 1 (1.75×10^{20} POT), run 2 (2.70×10^{20} POT), and run 3 (2.43×10^{20} POT) [148]. Data from two more MicroBooNE run periods exist on tape and are currently being processed by the collaboration, corresponding to another roughly 5×10^{20} POT. The entire dataset was taken during nominal neutrino mode running conditions for the BNB, in which 8 GeV kinetic energy protons from the Fermilab Booster strike the beryllium target to produce a cascade of mesons, and the positively-charged mesons are focused and decay to produce a beam of neutrinos. The BNB was described in detail in section 2.1.1. The intrinsic ν_e component of the beam, which makes up $\sim 0.5\%$ of the total neutrino flux, can be seen in figure 2-2.

The MicroBooNE simulation relies on the GENIE v3.00.06 [173–176] neutrino event generator. MicroBooNE specifically employs the G18_10a_02_11a model, which uses the Valencia CCQE and meson-exchange-current (MEC) models [177] as well as a local Fermi gas model to describe the nuclear environment. GENIE models the final state interactions (FSI) of particles produced in neutrino interactions with the nuclear environment using the typical $A^{2/3}$ scaling, including corrections from nuclear binding energy [178] as well as Pauli blocking and the velocity dependence of the nuclear mean field [179]. To improve the nuclear modeling of CCQE and MEC interactions, MicroBooNE has developed a custom tune of the G18_10a_02_11a

model using a ν_μ CC0 π cross section measurement from T2K [180,181]. This so-called “MicroBooNE tune” was determined by performing a two-dimensional fit of the T2K double-differential cross section data in muon momentum p_μ and angle $\cos\theta_\mu$. The fit adjusts four parameters related to the CCQE and MEC cross section models in GENIE v3.00.06, resulting in a better fit to both T2K and MiniBooNE double-differential cross section data (though discrepancies exist at lower muon momenta in the latter case). The MicroBooNE tune is used for all three MicroBooNE ν_e analyses; it gives better agreement with MicroBooNE ν_μ data, which are essential for constraining the ν_e prediction in each analysis, and updates the predicted ν_e CCQE and MEC cross sections on argon [180].

Electron scattering data can be used to evaluate the vector part of the vector-axial neutrino-nucleus cross section models in event generators such as GENIE v3.00.06 [182, 183]. Earlier versions of GENIE were found to be in disagreement with electron scattering data in the sub-GeV energy transfer regime [182]. Updates to the QE, MEC, and Δ excitation models in GENIE improved agreement with the QE peak and MEC-dominated dip region, though disagreements remain at higher momentum transfers [183]. Dedicated electron scattering measurements using the CLAS detector at Jefferson Lab have also revealed biases in the GENIE-based reconstruction of the incident electron energy, especially in the large final state transverse momentum regime [184].

Particle propagation in the liquid argon detector is handled by the Geant4 toolkit V10.3.03c [185], and the MicroBooNE TPC response is modeled using the general LArSoft [186] package and the internal uboonecode package. The effect of cosmic particles crossing the TPC during the readout window is modeled by overlaying off-beam data on top of the simulated events. This is important, as around 20-30 cosmic rays can cross the detector in a single event due to the long readout window and near-surface location of the MicroBooNE detector [148].

4.2 The Electron Low Energy Excess Template

The hypothesis studied in each of the three electron analyses is that the MiniBooNE anomaly comes from an excess of ν_e interactions from the BNB. While such an excess could come from $\nu_\mu \rightarrow \nu_e$ oscillations in a $3 + 1$ model, this is not the only possible source. For example, mismodeling of μ^+ decay-in-flight and K^+ decay-at-rest in the BNB [14] could lead to an enhancement of low energy ν_e compared with the expectation.

MicroBooNE has elected to remain agnostic about possible sources of ν_e events for its first results. To do this, the collaboration has developed an “electron low energy excess (eLEE) model” by unfolding the MiniBooNE excess under a ν_e hypothesis [20, 148]. This is accomplished using the D’Agostini iterative unfolding procedure [187], which updates a predicted distribution in true space using an observed distribution in reconstructed space and a known detector response matrix connecting the two. The algorithm relies on Bayes theorem to update the predicted distribution over a specified number of iterations. A small number of iterations will give an unfolded distribution biased toward the initial guess, while a larger number of iterations will give an unfolded distribution that is more sensitive to statistical fluctuations in the observed data [20]. For the eLEE model, the MiniBooNE observation as a function of reconstructed neutrino energy, as shown in figure 2-5, is used to unfold the predicted ν_e interaction rate in MiniBooNE as a function of the true neutrino energy. This procedure accounts for the detection efficiency; thus, the unfolded distribution can be interpreted as the number of ν_e events that interact in MiniBooNE. One can take the ratio of the unfolded MiniBooNE ν_e prediction to the nominal ν_e prediction from the MiniBooNE Monte Carlo (MC), as shown in figure 4-1a. This gives a set of binned weights as a function of the true ν_e energy, which can be applied to any BNB MC sample of ν_e interactions to represent a MiniBooNE-like excess. These weights constitute the MicroBooNE eLEE model of the MiniBooNE anomaly. Figure 4-1b shows the actual eLEE model weights used in all three MicroBooNE ν_e analyses, calculated using the first 12.84×10^{20} POT of the MiniBooNE neutrino mode dataset.

A few comments are necessary regarding the unfolding procedure. First, the number of iterations is a regularization parameter that must be selected by the user. MicroBooNE used three iterations of the D’Agostini algorithm to produce the eLEE model, a choice designed to minimize the variance and bias in the unfolded spectrum while remaining in good agreement with the MiniBooNE data, i.e. retaining a χ^2 per degree-of-freedom less than 1 in the reconstructed distribution. While it is true that these criteria select a unique choice for the number of iterations, the eLEE model shown in figure 4-1b is not the unique unfolded distribution consistent with the MiniBooNE excess. This point is discussed further in Ref. [188]. Additionally, the eLEE model was developed before MiniBooNE had looked at any ν_e data in the $E_\nu^{\text{QE}} < 200$ MeV region [79]. Thus, the unfolded ν_e prediction will lose events that reconstruct below 200 MeV. It is also worth noting that while the MiniBooNE e -like dataset is only sensitive in principle to ν_e CC interactions without pions in the final state (as both charged and neutral pions are visible in a Cherenkov detector), the eLEE model weights are applied to all ν_e CC interactions in MicroBooNE. The impact of this choice is mitigated by the development of three different ν_e analyses with an emphasis on different final states. If the MiniBooNE excess only exists in one class of ν_e interactions, this should be evident when comparing results from all three analyses. Finally, the eLEE model does not account for systematic uncertainties on the MiniBooNE prediction. This is important, as MiniBooNE is a systematics-limited experiment [10]. The eLEE model weights shown in figure 4-1b can be thought to represent a “median” model of the MiniBooNE excess. One can scale the weights by an overall factor, hereafter referred to as the “LEE signal strength” x_{LEE} , to capture the 21% statistical \oplus systematic normalization uncertainty on the MiniBooNE prediction. In chapter 5, we will cover results from a one-dimensional signal strength scaling test.

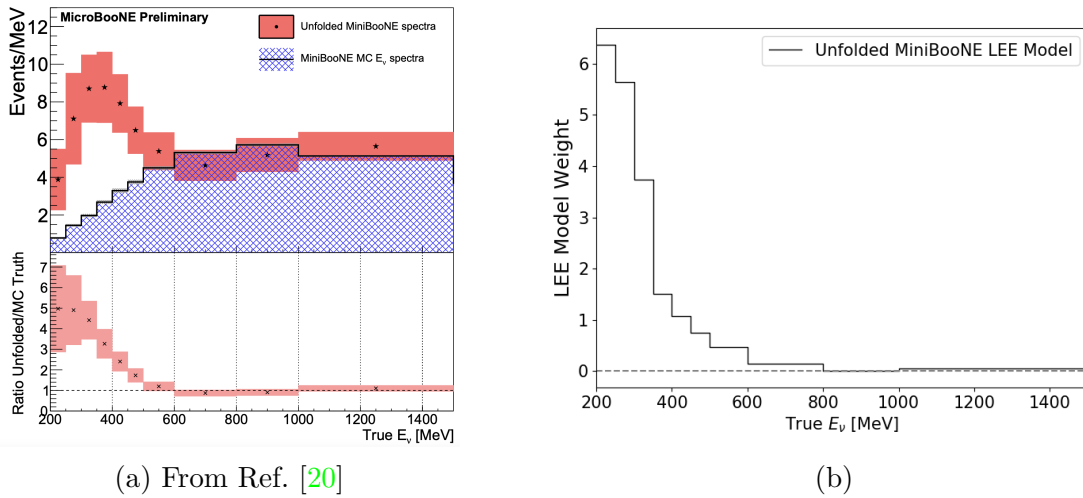


Figure 4-1: Figure 4-1a shows the unfolded ν_e prediction in MiniBooNE calculated using the first 6.46×10^{20} POT of neutrino mode data [20]. Figure 4-1b shows the unfolded eLEE model weights derived from the first 12.84×10^{20} POT of MiniBooNE neutrino mode data, which constitute the MicroBooNE eLEE model.

4.3 Philosophy Behind the Two-Body CCQE Analysis

The two-body CCQE analysis seeks to do what MiniBooNE could not— isolate a high-purity sample of ν_e interactions with a low SM photon background, especially at low energies. To do this, we have selected a very specific final state definition: ν_e CCQE interactions with reconstructed kinematic variables that are consistent with two-body scattering.

The ν_ℓ CCQE interaction channel is a simple two-body scattering process,

$$\nu_\ell n \rightarrow \ell p. \quad (4.1)$$

It was chosen because CCQE is the dominant neutrino interaction at lower neutrino energies, where the eLEE model weights in figure 4-1b are the highest. As shown in figure 4-2, CCQE interactions comprise 77% of the ν_e CC interactions in MicroBooNE below 500 MeV [148]. Other interaction channels, such as meson exchange current (MEC) and resonant (RES) scattering, become more significant at higher

neutrino energies. These channels are much more complicated than CCQE and thus come with larger modeling uncertainties. MEC events involve pion exchange between two nucleons and are thus sometimes referred to as “two-particle-two-hole” interactions [189]. RES events involve the excitation of a nucleon to a Δ resonance which will subsequently decay, most often to a pion and nucleon. CCQE events are also easier to reconstruct; the original neutrino energy is simply the sum of the observed electron and proton kinetic energies, up to smearing from nuclear effects. In MEC and RES events, one must account for the kinetic energy of additional final state protons and pions, and the potential existence of invisible final state neutrons makes it more difficult to connect the total kinetic energy of the observed final state to the initial neutrino energy. Thus, the relative simplicity of CCQE interactions is another advantage of choosing such a signal definition.

The two-body CCQE analysis has developed a number of techniques to isolate ν_e CCQE events from other ν_e and ν_μ interactions in MicroBooNE. The most basic of these involves leveraging the clean $1e1p$ final state topology of the signal events. The left panel of figure 4-3 shows an example candidate $1e1p$ event in MicroBooNE data. One can clearly see the proton track and electron EM shower emanating from the same vertex, presumably the location of the neutrino interaction in the LArTPC. Our reconstruction algorithm takes advantage of this distinct “vee” shape, imposing a strict requirement that exactly two “prongs” of charge (i.e., tracks or showers) come from the same 3D-consistent interaction vertex. This process is augmented by a powerful convolutional neural network to specifically find events with one track and one shower, indicative of a $1e1p$ final state. The reconstruction algorithm for this analysis will be discussed in more detail in the next section.

We also rely on two-body kinematics to select ν_e CCQE interactions. In theory, the observed $1e1p$ final state from a ν_e CCQE interaction should obey kinematic constraints imposed in a two-body scattering process. In a realistic MicroBooNE scenario, the struck neutron lives within the argon nuclear medium and has non-zero initial momentum. The final state proton must also travel through the argon nuclear medium, in which it can undergo final state interactions that alter its momentum

and/or create additional hadronic particles. In order to optimize our reconstruction and minimize uncertainties, we are interested in events with minimal final state interactions. We also want to remove backgrounds that will be especially inconsistent with two-body kinematics, such as ν_e MEC events and ν_μ events with a π^0 in the final state in which one of the photons is reconstructed alongside a proton track. This is achieved by introducing variables to the analysis that are especially sensitive to deviations from two-body kinematics. For example, the total transverse momentum of the neutrino interaction (p_T) will be further from zero for non-CCQE events. Additionally, as shown in figure 4-4, the proton in a two-body CCQE scattering interaction will always be forward-going. While nuclear effects can in principle break this requirement in MicroBooNE, it is much more likely for a non-CCQE background event to result in a backward-going proton than a true CCQE event. Therefore, we explicitly require forward-going protons in this analysis. We can also consider the Bjorken scaling variable

$$x_{Bj} = \frac{Q^2}{2p_p \cdot q}, \quad (4.2)$$

where $q^2 = -Q^2$ is the four-momentum transfer and p_p is the final state proton four-momentum. In the case of quasi-elastic scattering, x_{Bj} should be close to unity up to nuclear effects [190], while it can significantly differ from unity for non-CCQE backgrounds [148].

Finally, we can harness the full power of the LArTPC technology to reconstruct the ν_e energy in multiple ways. For ν_e traveling along the z axis of the detector, the neutrino energy can be calculated from any of the following,

$$E_\nu^{\text{range}} = K_p + K_\ell + M_\ell + M_p - (M_n - B), \quad (4.3)$$

$$E_\nu^{QE-p} = \left(\frac{1}{2}\right) \frac{2 \cdot (M_n - B) \cdot E_p - ((M_n - B)^2 + M_p^2 - M_\ell^2)}{(M_n - B) - E_p + \sqrt{(E_p^2 - M_p^2)} \cdot \cos \theta_p}, \quad (4.4)$$

$$E_\nu^{QE-\ell} = \left(\frac{1}{2}\right) \frac{2 \cdot (M_n - B) \cdot E_\ell - ((M_n - B)^2 + M_\ell^2 - M_p^2)}{(M_n - B) - E_\ell + \sqrt{(E_\ell^2 - M_\ell^2)} \cdot \cos \theta_\ell}, \quad (4.5)$$

where $K_{\ell/p}$ is the kinetic energy of the lepton or proton, $\theta_{\ell/p}$ is the angle of the

shower or track with respect to the z axis, $M_{\ell/p/n}$ is the particle mass, and B is the average binding energy, which is around 40 MeV in argon [191]. For a true ν_e CCQE interaction, all three of these formulae should yield a similar reconstructed neutrino energy. To leverage this effect we define the ‘‘QE consistency’’ variable Δ^{QE} as

$$(\Delta^{QE})^2 \equiv (E_\nu^{\text{range}} - E_\nu^{QE-p})^2 + (E_\nu^{\text{range}} - E_\nu^{QE-\ell})^2 + (E_\nu^{QE-\ell} - E_\nu^{QE-p})^2, \quad (4.6)$$

which should be close to zero for CCQE events. To mitigate smearing from nuclear effects, E_ν^{QE-p} and $E_\nu^{QE-\ell}$ are calculated in the rest frame of the struck nucleon. We note here that equation (4.3) is the official reconstructed neutrino energy used for this analysis; E_ν and E_ν^{range} are used interchangeably throughout this thesis.

Finally, it is worth mentioning that the two-body CCQE analysis also isolates a control sample of ν_μ CCQE interactions with a $1\mu 1p$ final state. The $1\mu 1p$ sample is constructed analogously to the signal $1e 1p$ sample, though with around two orders of magnitude more events (as ν_μ comprise most of the BNB flux). This sample is used to constrain the intrinsic ν_e prediction and uncertainties in the $1e 1p$ sample, as the ν_e and ν_μ flux and cross section are highly correlated. We also make use of dedicated π^0 and Michel electron samples, as will be discussed in section 4.4.3.

In summary, the goal of the analysis presented here is the isolation of clean two-body ν_e CCQE interactions in MicroBooNE. The following sections explain how we leverage the topological and kinematic patterns described above to identify such events.

4.4 Reconstruction

We begin with a broad overview of the reconstruction chain for the two-body CCQE analysis. An initial set of neutrino event candidates are identified by requiring sufficient light detection in time with the beam trigger, as discussed in section 3.1.3. These events are passed through the noise filter and deconvolution algorithm described in section 3.2 and shown in figure 3-7. Additionally, a ‘‘good runs cut’’ is




Figure 4-2: The expected distribution of ν_e interactions in MicroBooNE as a function of the true ν_e energy. The dotted line shows the expectation from the MicroBooNE eLEE model of the MiniBooNE excess discussed in section 4.2.




Figure 4-3: An example candidate $1e1p$ data event in MicroBooNE, including the raw LArTPC collection plane image (left) and the pixel labels assigned from SparseSSNet (right).

employed to remove events during which the MicroBooNE detector response was especially noisy [148]. These steps are common to all MicroBooNE analyses. They result in one event display per wire plane for every event, similar to the rightmost panel of figure 3-7 and the left panel of figure 4-3. The color axis in the left panel of figure 4-3 is given in “pixel intensity units” (PIU), which are defined as the integrated charge in one wire over six $0.5 \mu\text{s}$ time ticks [148]. Given the electron drift speed of $0.11 \text{ cm}/\mu\text{s}$, this corresponds to a spatial extent of 0.33 cm along the drift direction. The wire spacing is such that each pixel in the left panel of figure 4-3 represents a $0.3 \text{ cm} \times 0.33 \text{ cm}$ 2D square in the detector.




Figure 4-4: The relationship between the final state lepton (left) and proton (right) energy and scattering angle, for different neutrino energies, in ν_μ (top) and ν_e (bottom) CCQE scattering.

In the two-body CCQE analysis, these event displays are first passed through the sparse semantic segmentation network (**SparseSSNet**) [21]. **SparseSSNet** is a custom convolutional neural network (CNN) developed by the analysis team that labels the active pixels in each event according to five different charge deposition hypotheses: highly-ionizing particle (HIP), minimum-ionizing particle (MIP), EM activity, delta electron, or Michel electron. For the purposes of this analysis, we group HIP and MIP labels into the parent “track” label and EM activity, delta electrons, and Michel electrons into the parent “shower” label. An example **SparseSSNet** label is shown in the right panel of figure 4-3.

After going through **SparseSSNet**, the images are filtered for crossing cosmic interactions using the WireCell charge-to-light matching algorithm [192, 193]. This algorithm matches charge observed on the TPC to light observed by the PMTs to identify regions of charge which are not coincident in time with the beam spill, presumably coming from cosmic rays. We can then mask these pixels to remove the bulk of cosmic-related charge deposition. The cosmic-filtered **SparseSSNet** output is then used as input to a vertex and track identification algorithm, which looks for a shower-

like and track-like prong of charge coming from a given 3D vertex [24]. Track-like prongs pass through a dedicated step of the algorithm that performs a step-by-step 3D reconstruction of the extent and direction of the track, which is then translated into a range-based energy and angle estimation. The events are then passed through a dedicated EM shower identification algorithm that identifies cones of shower-like pixels pointing back to the neutrino interaction vertex and estimates the total visible energy and angle of the shower. [172]. This algorithm has the capability of finding up to two showers separated from the vertex, which is necessary for identifying events with a $\pi^0 \rightarrow \gamma\gamma$ decay.

After finding candidate neutrino vertices, we pass the event displays through the multiple particle identification (MPID) CNN [22]. This network assigns five scores to each image indicating the network’s confidence that the event contains a proton, electron, muon, photon, or charged pion. The MPID scores are not used until the final signal selection described in section 4.5.

At this point, the reconstruction chain has identified two-prong candidate events and reconstructed the relevant lowest-level kinematic variables of each final state particle, i.e. the energy and direction. These kinematic variables are then combined to calculate more complicated kinematic variables such as Q^2 , p_T , x_{Bj} , and the reconstructed π^0 mass (in the case of two-shower events), as well as the most important variable in this analysis: the reconstructed neutrino energy E_ν^{range} , as defined in equation (4.3). The full suite of reconstructed kinematic variables used to define the $1e1p$ and $1\mu1p$ samples is given in table 4.1. These variables are used as input to the signal selection described in section 4.5. The following subsections explain in more detail the most important steps of the reconstruction procedure.

4.4.1 Convolutional Neural Networks in LArTPCs

As discussed above, The two-body CCQE analysis relies on two CNNs to assist in reconstructing the MicroBooNE LArTPC images: `SparseSSNet` and `MPID`.

`SparseSSNet` is described in detail in Ref. [21]. The network uses a combination of the U-Net [194] and ResNet [195] architectures. The U-Net uses an encoding step,

in which the input image is downsampled to a smaller size, followed by a decoding step, in which the encoded image is upsampled to the same dimensionality and size as the original image. This second step is required to perform pixel-level labeling. The ResNet architecture employs residual connections that improve the training performance of deeper networks compared to regular CNNs [195]. **SparseSSNet** also makes use of sparse submanifold convolutions [196], which perform better on semantic segmentation tasks for sparse input data like LArTPC images (where very few pixels in a given image have nonzero charge). The training sample is generated using the simulation described in section 4.1; specifically, a random set of visible final state particles are generated at randomly distributed locations within the detector. The network performs considerably well on the nominal MicroBooNE BNB simulation sample, as well as the dedicated intrinsic ν_e simulation sample, neither of which appeared in the **SparseSSNet** training sample [21]. While **SparseSSNet** has not been statistically evaluated against MicroBooNE data, hand scans of the network output on data event displays do not reveal any bias in the network performance [21]. The performance of an earlier iteration of **SparseSSNet** was evaluated using two data samples: Michel electrons and ν_μ CC π^0 events [197]. Good agreement was observed between pixel labels in data and simulation for both samples, and no significant differences were found between network-labeled images and physicist-labeled images. Figure 4-5 shows the **SparseSSNet** architecture as well as an example labeled image in the training dataset.

The MPID network is described in detail in Ref. [22]. It uses a standard CNN architecture which takes as input a 2D LArTPC image cropped around the neutrino interaction candidate vertex on a single wire plane and outputs five numbers representing different particle scores for protons, electrons, muons, photons, and charged pions. The score is meant to represent the likelihood that a given particle exists in the provided image. Two types of scores are generated: an “image” score, describing the probability that the particle exists anywhere in the image, and the “interaction” score, describing the probability that the particle is connected to the neutrino interaction vertex. The training sample is generated similarly to that of **SparseSSNet**; random sets of particles are generated throughout the detector volume with energies

in a range relevant to MicroBooNE. The MPID score distributions exhibit excellent agreement between data and simulation in two different samples: ν_μ CC $1\mu 1p$ events and ν_μ CC π^0 events [22]. Figure 4-6 shows the MPID architecture as well as an example simulated image with MPID image scores.

It is also worth briefly mentioning energy reconstruction in LArTPCs using CNNs. My earliest work as a graduate student investigated the ability of the ResNet [195] and Inception network [198] CNN models to reconstruct EM shower energies in MicroBooNE. Compared to the standard clustering-based shower reconstruction algorithm described in section 4.4.3, the CNN method appeared to better handle EM showers that passed through unresponsive wires in the MicroBooNE detector. Despite the promise of the CNN method, it was not developed in time to be implemented in the two-body CCQE analysis. Nevertheless, the MicroBooNE-specific network was extended by an MIT undergraduate, Kiara Carloni, to enable a systematic study of CNN-based EM shower energy reconstruction in a general LArTPC [23]. The corresponding publication is included in Appendix A. Figure 4-7 shows the main takeaway of this work: a CNN-based reconstruction method is significantly more robust to unresponsive wires in the detector compared with a traditional EM shower reconstruction algorithm, which employs a linear calibration between the total shower charge and energy.

4.4.2 Vertex and Track Reconstruction

After running `SparseSSNet` and filtering cosmic-associated charge, the images are passed through the vertex reconstruction algorithm described in Ref. [24]. In brief, the algorithm looks for the characteristic “vee” shape of both $1e1p$ signal events, as shown in the left panel of figure 4-3, and $1\mu 1p$ control sample events. The algorithm begins by isolating a sample of 2D vertex seeds (i.e., candidate neutrino interaction locations) on each plane. This is done by clustering regions of charge and then applying two techniques—contour defect identification using a convex hull and intersection identification using principal component analysis. The set of vertex seeds is reduced to those that are consistent in the time dimension across all three planes. From each

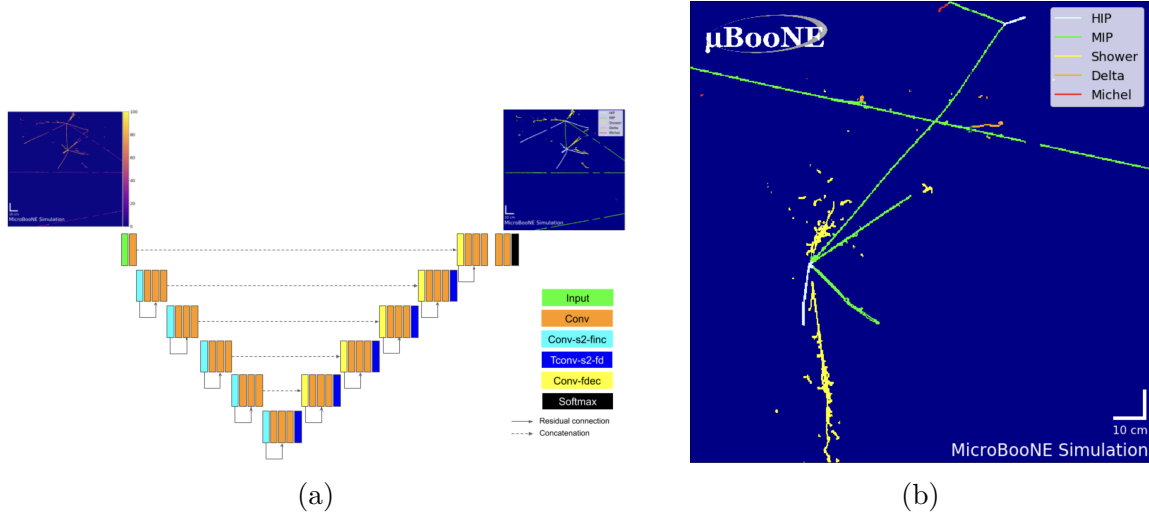
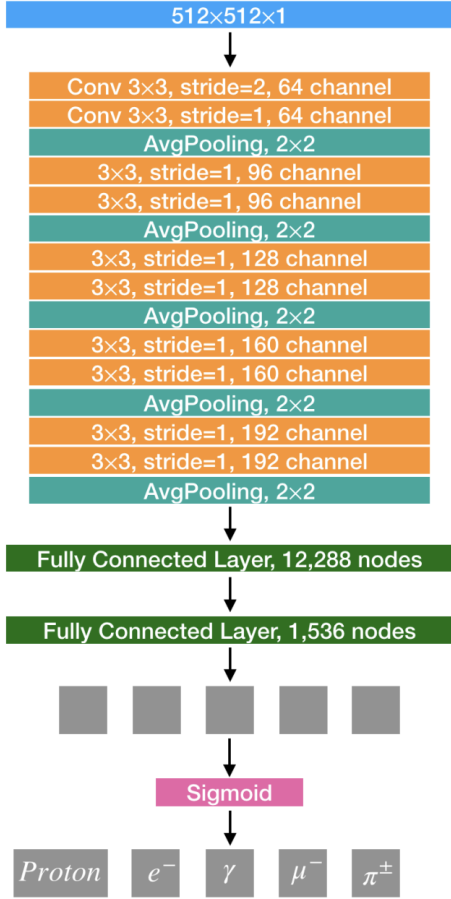


Figure 4-5: Figure 4-5a shows a diagram of the SparseSSNet U-ResNet architecture. Figure 4-5b shows the SparseSSNet pixel labels on a simulated image in the training dataset. Figures from Ref. [21].


triplet of time-consistent vertex seeds on the three wire planes, a final 3D vertex is found by minimizing an angular metric that represents the likelihood that two clusters of charge emanate radially from that vertex location. Figure 4-8a shows an example minimization of this angular metric, which involves calculating two different estimates of the opening angle between the two prongs for a given vertex location. Simulation studies indicate that this algorithm is able to reconstruct neutrino interaction vertices to within less than 1 cm with an average efficiency of 56% [24].

After finding a neutrino vertex candidate, we perform a 3D reconstruction of the prongs of charge emitted from the vertex [24]. An iterative stochastic algorithm identifies a set of 3D spatial points associated with each prong by building upon previous points; these points comprise the track. This process is shown diagrammatically in figure 4-8b. It is worth noting that the “tracks” at this stage can correspond to either true tracks (HIPs and MIPs) or EM showers.

The 3D track can be projected onto each 2D plane. At this point, the SparseSSNet labels on each plane are used to determine if the event is a “track-track” or “track-shower” event. The latter case happens when greater than 20% of the pixels in a given prong are labeled as shower pixels. In the “track-shower” case, the shower-like prong is considered to be the electron while the track-like prong is the proton. Thus, these



(a)



	p	e^-	γ	μ^-	π^\pm
MPID Score	0.89	0.95	0.85	0.06	0.17

(b)

Figure 4-6: Figure 4-6a shows a diagram of the MPID architecture. Figure 4-6b shows the MPID image scores on an example simulated event. Figures from Ref. [22].

are candidate $1e1p$ events. The “track-track” case corresponds to candidate $1\mu1p$ event, where the shorter track is ascribed to the highly-ionizing proton. Electron EM showers are not sufficiently reconstructed in this step and are left for a dedicated shower reconstruction algorithm described in section 4.4.3. Muons and protons can, however, be reconstructed at this stage. The polar and azimuth angles of each particle can be determined from the direction of the associated track. The length of the tracks is used in conjunction with the known stopping power of each particle in liquid argon to determine the kinetic energy. Simulation studies of $1\mu1p$ events suggest an angular resolution of a few degrees and a kinetic energy resolution of a few percent [24].

It is also worth mentioning that two-prong events with an opening angle of less




Figure 4-7: The fraction of EM showers reconstructed to within 5% of the true deposited energy as a function of the fraction of unresponsive wires in the LArTPC collection plane, considering three different methods. The ResNet and Inception network significantly outperform the traditional linear calibration between charge and energy. Figure from Ref. [23].

than 10° in any plane are rejected at this point. Thus, signal events in this analysis consisted of well-separated leptons and protons.

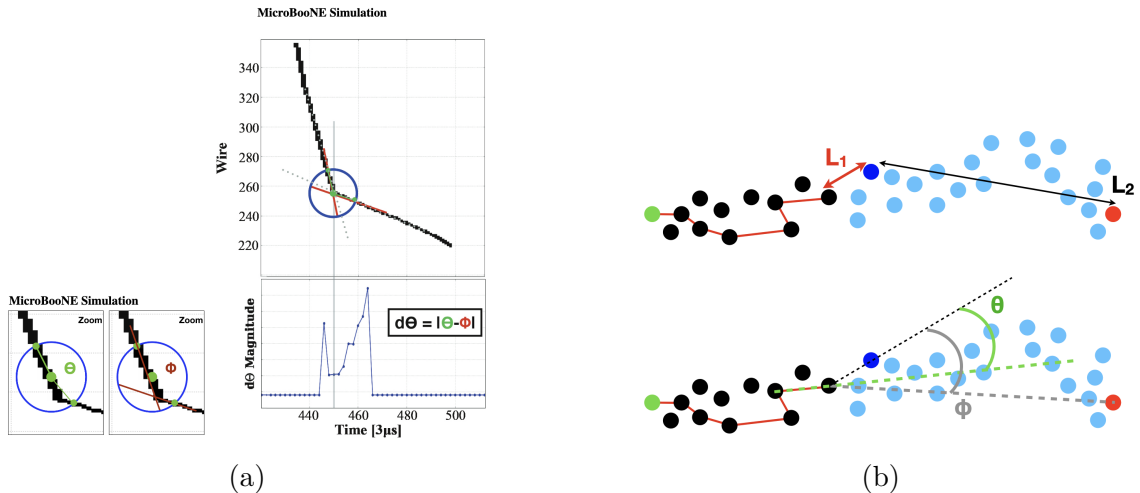


Figure 4-8: Figure 4-8a shows the angular metric that is minimized to find a 3D neutrino vertex candidate. Figure 4-8b shows the iterative track reconstruction algorithm, which relies on calculating distances (L_1 and L_2) and angles (θ and ϕ) with respect to the previous point and the end of the track. Figures from Ref. [24].

4.4.3 Publication: *Electromagnetic shower reconstruction and energy validation with Michel electrons and π^0 samples for the deep-learning-based analyses in MicroBooNE*

As mentioned in the previous section, a dedicated algorithm is used to reconstruct the EM showers coming from the neutrino interaction vertex. The algorithm is described in detail in Ref. [172], which also covers a data-driven validation of the shower reconstruction using dedicated π^0 and Michel electron samples. A brief overview of the shower reconstruction is as follows. We first mask out pixels in each plane for which the `SparseSSNet` shower score is below a value of 0.5. Next, we attach a triangle to the interaction vertex in each plane and optimize the direction, gap from the vertex, opening angle, and length of the triangle to encapsulate as many shower pixels as possible without becoming too large. We calculate the energy of the shower by performing a linear calibration between the summed PIU in the collection plane shower triangle and the simulated electron/photon energy. The collection plane is used because it produces the most robust signals, as discussed in section 3.2. We reconstruct 3D showers by looking for overlap in the time dimension between the shower triangles on all three planes. This algorithm can be run a second time after masking out the optimized triangle from the first pass, which is essential for isolating a π^0 sample.

This MicroBooNE publication was led by myself and Katie Mason, a graduate student from Tufts University. Katie focused on the triangle-fitting procedure of the shower reconstruction and the π^0 mass peak fit, while I focused on the charge-to-energy scaling of the shower reconstruction and the Michel energy spectrum fit. The full *JINST* publication is included below. The most important result from this study is shown in Figure 16 of the paper, which demonstrates good agreement between the simulation-derived charge-to-energy scaling applied to EM showers in this analysis and two standard candles: the π^0 invariant mass peak and the cutoff of the Michel electron energy distribution.

TECHNICAL REPORT

Electromagnetic shower reconstruction and energy validation with Michel electrons and π^0 samples for the deep-learning-based analyses in MicroBooNE

To cite this article: The MicroBooNE Collaboration *et al* 2021 *JINST* **16** T12017

View the [article online](#) for updates and enhancements.

You may also like

- [A muon-track reconstruction exploiting stochastic losses for large-scale Cherenkov detectors](#)
on behalf of the IceCube collaboration, R. Abbasi, M. Ackermann et al.
- [Comparative evaluation of analogue front-end designs for the CMS Inner Tracker at the High Luminosity LHC](#)
The Tracker Group of the CMS Collaboration, W. Adam, T. Bergauer et al.
- [Measurement of the atmospheric muon rate with the MicroBooNE Liquid Argon TPC](#)
MicroBooNE, P. Abratenko, M. Alrashed et al.

TECHNICAL REPORT

Electromagnetic shower reconstruction and energy validation with Michel electrons and π^0 samples for the deep-learning-based analyses in MicroBooNE

The MicroBooNE Collaboration

P. Abratenko,^{gg} R. An,ⁿ J. Anthony,^d L. Arellano,^r J. Asaadi,^{ff} A. Ashkenazi,^{dd} S. Balasubramanian,^k B. Baller,^k C. Barnes,^t G. Barr,^w V. Basque,^r L. Bathe-Peters,^m O. Benevides Rodrigues,^{cc} S. Berkman,^k A. Bhanderi,^r A. Bhat,^{cc} M. Bishai,^b A. Blake,^p T. Bolton,^o J.Y. Book,^m L. Camilleri,ⁱ D. Caratelli,^k I. Caro Terrazas,^h R. Castillo Fernandez,^k F. Cavanna,^k G. Cerati,^k Y. Chen,^a D. Cianci,ⁱ J.M. Conrad,^s M. Convery,^z L. Cooper-Troendle,^{jj} J.I. Crespo-Anadón,^e M. Del Tutto,^k S.R. Dennis,^d P. Detje,^d A. Devitt,^p R. Diurba,^u R. Dorrill,ⁿ K. Duffy,^k S. Dytman,^x B. Eberly,^{bb} A. Ereditato,^a J.J. Evans,^r R. Fine,^q G.A. Fiorentini Aguirre,^{aa} R.S. Fitzpatrick,^t B.T. Fleming,^{jj} N. Foppiani,^m D. Franco,^{jj} A.P. Furmanski,^u D. Garcia-Gamez,^l S. Gardiner,^k G. Ge,ⁱ S. Gollapinni,^{ee,q} O. Goodwin,^r E. Gramellini,^k P. Green,^r H. Greenlee,^k W. Gu,^b R. Guenette,^m P. Guzowski,^r L. Hagaman,^{jj} O. Hen,^s C. Hilgenberg,^u G.A. Horton-Smith,^o A. Hourlier,^s R. Itay,^z C. James,^k X. Ji,^b L. Jiang,^{hh} J.H. Jo,^{jj} R.A. Johnson,^g Y.-J. Jwa,ⁱ D. Kalra,ⁱ N. Kamp,^s N. Kaneshige,^c G. Karagiorgi,ⁱ W. Ketchum,^k M. Kirby,^k T. Kobilarcik,^k I. Kreslo,^a R. LaZur,^h I. Lepetic,^y K. Li,^{jj} Y. Li,^b K. Lin,^q B.R. Littlejohn,ⁿ W.C. Louis,^q X. Luo,^c K. Manivannan,^{cc} C. Mariani,^{hh} D. Marsden,^r J. Marshall,ⁱⁱ D.A. Martinez Caicedo,^{aa} K. Mason,^{gg} A. Mastbaum,^y N. McConkey,^r V. Meddage,^o T. Mettler,^a K. Miller,^f J. Mills,^{gg} K. Mistry,^r A. Mogan,^{ee} T. Mohayai,^k J. Moon,^s M. Mooney,^h A.F. Moor,^d C.D. Moore,^k L. Mora Lepin,^r J. Mousseau,^t M. Murphy,^{hh} D. Naples,^x A. Navrer-Agasson,^r M. Nebot-Guinot,^j R.K. Neely,^o D.A. Newmark,^q J. Nowak,^p M. Nunes,^{cc} O. Palamara,^k V. Paolone,^x A. Papadopoulou,^s V. Papavassiliou,^v S.F. Pate,^v N. Patel,^p A. Paudel,^o Z. Pavlovic,^k E. Piasetzky,^{dd} I.D. Ponce-Pinto,^{jj} S. Prince,^m X. Qian,^b J.L. Raaf,^k V. Radeka,^b A. Rafique,^o M. Reggiani-Guzzo,^r L. Ren,^v L.C.J. Rice,^x L. Rochester,^z J. Rodriguez Rondon,^{aa} M. Rosenberg,^x M. Ross-Lonergan,ⁱ G. Scanavini,^{jj} D.W. Schmitz,^f A. Schukraft,^k W. Seligman,ⁱ M.H. Shaevitz,ⁱ R. Sharankova,^{gg} J. Shi,^d J. Sinclair,^a A. Smith,^d E.L. Snider,^k M. Soderberg,^{cc} S. Söldner-Rembold,^r P. Spentzouris,^k J. Spitz,^t M. Stancari,^k J. St.,^k T. Strauss,^k K. Sutton,ⁱ S. Sword-Fehlberg,^v A.M. Szecel,^j N. Tagg,^w W. Tang,^{ee} K. Terao,^z C. Thorpe,^p D. Totani,^c M. Toups,^k Y.-T. Tsai,^z M.A. Uchida,^d T. Usher,^z W. Van De Pontseele,^{w,m} B. Viren,^b M. Weber,^a H. Wei,^b Z. Williams,^{ff} S. Wolbers,^k T. Wongjirad,^{gg} M. Wospakrik,^k K. Wresilo,^d N. Wright,^s W. Wu,^k E. Yandel,^c T. Yang,^k G. Yarbrough,^{ee} L.E. Yates,^s H.W. Yu,^b G.P. Zeller,^k J. Zennamo,^k and C. Zhang^b

- ^a *Universität Bern, Bern CH-3012, Switzerland*
- ^b *Brookhaven National Laboratory (BNL), Upton, NY, 11973, U.S.A.*
- ^c *University of California, Santa Barbara, CA, 93106, U.S.A.*
- ^d *University of Cambridge, Cambridge CB3 0HE, United Kingdom*
- ^e *Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid E-28040, Spain*
- ^f *University of Chicago, Chicago, IL, 60637, U.S.A.*
- ^g *University of Cincinnati, Cincinnati, OH, 45221, U.S.A.*
- ^h *Colorado State University, Fort Collins, CO, 80523, U.S.A.*
- ⁱ *Columbia University, New York, NY, 10027, U.S.A.*
- ^j *University of Edinburgh, Edinburgh EH9 3FD, United Kingdom*
- ^k *Fermi National Accelerator Laboratory (FNAL), Batavia, IL 60510, U.S.A.*
- ^l *Universidad de Granada, E-18071, Granada, Spain*
- ^m *Harvard University, Cambridge, MA 02138, U.S.A.*
- ⁿ *Illinois Institute of Technology (IIT), Chicago, IL 60616, U.S.A.*
- ^o *Kansas State University (KSU), Manhattan, KS, 66506, U.S.A.*
- ^p *Lancaster University, Lancaster LA1 4YW, United Kingdom*
- ^q *Los Alamos National Laboratory (LANL), Los Alamos, NM, 87545, U.S.A.*
- ^r *The University of Manchester, Manchester M13 9PL, United Kingdom*
- ^s *Massachusetts Institute of Technology (MIT), Cambridge, MA, 02139, U.S.A.*
- ^t *University of Michigan, Ann Arbor, MI, 48109, U.S.A.*
- ^u *University of Minnesota, Minneapolis, Mn, 55455, U.S.A.*
- ^v *New Mexico State University (NMSU), Las Cruces, NM, 88003, U.S.A.*
- ^w *University of Oxford, Oxford OX1 3RH, United Kingdom*
- ^x *University of Pittsburgh, Pittsburgh, PA, 15260, U.S.A.*
- ^y *Rutgers University, Piscataway, NJ, 08854, U.S.A., PA*
- ^z *SLAC National Accelerator Laboratory, Menlo Park, CA, 94025, U.S.A.*
- ^{aa} *South Dakota School of Mines and Technology (SDSMT), Rapid City, SD, 57701, U.S.A.*
- ^{bb} *University of Southern Maine, Portland, ME, 04104, U.S.A.*
- ^{cc} *Syracuse University, Syracuse, NY, 13244, U.S.A.*
- ^{dd} *Tel Aviv University, Tel Aviv, Israel, 69978*
- ^{ee} *University of Tennessee, Knoxville, TN, 37996, U.S.A.*
- ^{ff} *University of Texas, Arlington, TX, 76019, U.S.A.*
- ^{gg} *Tufts University, Medford, MA, 02155, U.S.A.*
- ^{hh} *Center for Neutrino Physics, Virginia Tech, Blacksburg, VA, 24061, U.S.A.*
- ⁱⁱ *University of Warwick, Coventry CV4 7AL, United Kingdom*
- ^{jj} *Wright Laboratory, Department of Physics, Yale University, New Haven, CT, 06520, U.S.A.*

E-mail: microboone_info@fnal.gov, John

ABSTRACT: This article presents the reconstruction of the electromagnetic activity from electrons and photons (showers) used in the MicroBooNE deep learning-based low energy electron search. The reconstruction algorithm uses a combination of traditional and deep learning-based techniques to estimate shower energies. We validate these predictions using two ν_μ -sourced data samples: charged/neutral current interactions with final state neutral pions and charged current interactions

in which the muon stops and decays within the detector producing a Michel electron. Both the neutral pion sample and Michel electron sample demonstrate agreement between data and simulation. Further, the absolute shower energy scale is shown to be consistent with the relevant physical constant of each sample: the neutral pion mass peak and the Michel energy cutoff.

KEYWORDS: Neutrino detectors; Noble liquid detectors (scintillation, ionization, double-phase); Pattern recognition, cluster finding, calibration and fitting methods; Time projection Chambers (TPC)

ARXIV EPRINT: [2110.11874](https://arxiv.org/abs/2110.11874)

2021 JINST 16 T12017

Contents

1	Introduction	1
2	Preliminary reconstruction of events in the deep learning analysis	2
3	Electromagnetic shower energy reconstruction	4
4	Neutral pion sample	8
4.1	Identification and reconstruction	8
4.2	Validation of agreement between simulation, data, and true rest mass	13
5	Michel electron sample	16
5.1	Identification and reconstruction	16
5.2	Validation of agreement between simulation, data, and true Michel cutoff	19
6	Combined validation of reconstructed shower energy	22
7	Conclusions	23

1 Introduction

The primary goal of the MicroBooNE experiment is investigate the anomalous excess of electron-like events observed in the MiniBooNE detector [1]. The anomaly is an excess of single electromagnetic showers with a peak shower energy in the 200-500 MeV observed in the MiniBooNE Cherenkov detector. The anomaly is here referred to as the Low Energy Excess (LEE). The MicroBooNE collaboration has developed several analyses designed to isolate LEE events, where event here refers to one detector readout record. The reconstruction tools presented here make up the shower reconstruction phase of the deep learning (DL)-based analysis [2]. A preliminary version of the (DL)-based analysis utilized in this study can be seen in ref. [3]. This analysis isolates events with 1 electron and 1 proton ($1e1p$) in the final state. The neutrino energy is reconstructed as:

$$E_\nu = K_p + K_e + M_e + M_p - (M_n - B), \quad (1.1)$$

where K indicates kinetic energy, M is mass, B is nuclear binding energy, and e and p indices indicate the electron and proton, respectively. The proton kinetic energy is reconstructed from the length of the track and the known energy deposited per unit length in liquid argon [4]. This article describes the reconstruction of the electron kinetic energy, which, for our ν_e signal of interest, ranges from 35 MeV to 1200 MeV. Across this range, the topology of the deposited electron energy changes from track-like at low energy to shower-like at high energy where photon radiation dominates. In this article, all electron energy deposits will be called showers despite the variety of topologies. The discussion will include both electrons and photons.

This article reports the method of electromagnetic shower reconstruction used in the DL-based analyses. The ν_e simulation-derived shower-charge-to-energy conversion value is presented in section 3. We then use two low-energy samples to validate the shower reconstruction. The first sample consists of photons produced by neutral pions (π^0) decays and is described in section 4. The second sample consists of Michel electrons produced when a stopped muon from a ν_μ charged current (CC) interaction decays and is described in section 5. We provide data-simulation comparison plots for each sample that we use to validate the relative shower energy scale between data and simulation. The use of the ν_e simulation-derived energy calculation is then further verified on data by measuring the agreement of data and simulation with two well-measured physical quantities: the π^0 invariant mass and the cutoff of the Michel energy spectrum.

2 Preliminary reconstruction of events in the deep learning analysis

The MicroBooNE detector is a 2.6 m \times 2.3 m \times 10.4 m Liquid Argon (LAr) Time Projection Chamber (TPC) filled with 85 metric tons of LAr [5]. The ionization electrons produced by charged particles in the event drift with a velocity of 0.1098 cm/ μ s through an applied potential of -70 kV to three wire planes [6]. The orientation of wires in the induction planes, U and V , are $+60^\circ$ and -60° relative to vertical. The collection plane, Y , has vertical wires. The wire spacing in all planes is 0.3 cm. The detector also contains a light collection system made up of 32 photomultiplier tubes that are used for triggering and initial event selection. The light collection system is not leveraged in the shower reconstruction work presented here.

The studies in this article use MicroBooNE data taken from 2016 to 2018 which were recorded over three run periods with 1.75×10^{20} protons-on-target (POT) in Run 1, 2.70×10^{20} POT in Run 2 and 2.43×10^{20} POT in Run 3. There are various Monte Carlo (MC) simulation samples that are used to build the reconstruction algorithm. The simulated neutrino events are overlaid with off-beam cosmic muon data. These “overlay samples” are used throughout the rest of this article. They contain all types of simulated neutrino events expected in the given POT. There are three samples of this type, one corresponding to cosmic information collected in each of the three MicroBooNE data runs used in the study which corresponds to 6.67×10^{20} total POT. Additionally, for the π^0 study we incorporate a specialized overlay sample containing a larger number of events with a π^0 that decays to two photons in the final state (high POT π^0 sample). The incorporation of this sample allows for a reduction of the statistical uncertainty on the simulation. For the Michel study, we incorporate an overlay sample of neutrino events interacting outside of the MicroBooNE detector, as muons from these external ν_μ CC interactions which enter the detector, come to a stop, and then decay provide an additional source of Michel electrons.

Various final state topologies occur at the neutrino energies observed in MicroBooNE. Those relevant to the DL-based LEE analysis have two particles attached to a vertex: $1e1p$ and $1\mu1p$ events from ν_e and ν_μ charged-current quasi-elastic or meson exchange current scattering, other types of ν_μ induced events with only two particles reconstructed at the vertex (such as π^0 events with disconnected photons), and cosmic ray muons. The first steps of the “low level reconstruction” is to isolate the two-particle-vertices of interest. The methods have been described elsewhere [3, 4]; we briefly review them here.

Preparation for reconstruction begins with algorithms to tag and discard the charge associated with the cosmic rays [7, 8]. Next the waveform data from the wires in each plane are converted to “images” which are two-dimensional distributions with wire number along the x axis and drift time along the y axis. The intensity of each image “pixel” is given by the integrated reconstructed charge waveform from six $0.5\ \mu\text{s}$ TPC time increments after applying noise-filtering [9] and signal processing [10, 11]. The six $0.5\ \mu\text{s}$ TPC time increments is comparable to the 3 mm wire spacing after accounting for the electron drift velocity. The intensity of each pixel is referred to as “ Q (charge)” in this article. An example event display of this data format is shown in figure 1 (a). The event display shown is of the collection plane (Y plane) of a simulated CC π^0 event which passes the shower reconstruction stage. The z -axis represents the charge (Q) at each (wire,time) pixel. To make the deposited charge more clear in the display, Q has been given a maximum value of 100. In the event shown, the image is cropped centering at the neutrino interaction vertex. Shower reconstruction is also shown which will be discussed in section 3.

Any charge tagged as associated with a cosmic ray is removed at this stage. The U , V , and Y images are then passed into the deep learning convolutional neural net, called “SparseSSnet,” which is a semantic segmentation algorithm that labels pixels as track-like or shower-like [12]. Figure 1 (b) shows the SparseSSNet shower scores of the same event as (a) with the track-like particles masked out.

The next step is to identify a “two particle” vertex, followed by 3D reconstruction of each particle. This process is described in detail for the $1\mu 1p$ sample in ref. [4], and the $1e 1p$ reconstruction follows similar steps. In short, the vertex algorithm searches for a characteristic “vee” shape where two particles meet at a vertex identifying cases where the particles are longer than 3 cm and the opening angle in at least one plane is greater than 10° . Based on the SparseSSnet pixel tagging, the vee may be formed of a “track-track” pair or a “shower-track” pair [12]. For any given vertex, more than one vee may be found if there are more than two particles emitted from the interaction point. Also, more than one vertex may be found at different points in the same event. An important example of this, used in this article, is the case of $1\mu 1p$ final state in which the muon decays to a Michel electron. This results in one vertex at the interaction point and another at the decay point. Because vertices are also found on cosmic rays, particularly those that stop and decay, many vertices are identified in any given event. All are passed to 3D reconstruction. At this stage in the reconstruction, multiple vertices can be reconstructed in each event. In later selections, a single vertex from the event will be chosen.

3D reconstruction of track-like objects (primarily protons and muons) is described in detail in ref. [4] and summarized here. The algorithm begins at the reconstructed vertex and follows ionization trails outward in 3D, clustering the charge into “prongs.” A prong is defined here as a collection of continuous charge, but it need not be a single line of charge. The prong may comprise connected branches of charge as will be the case for electromagnetic showers. Each prong is assumed to come from one particle. In the case where two prongs are identified, we identify the proton as the prong having the higher average pixel-based ionization density.

A fiducial volume containment requirement is enforced on all prongs. This requirement uses the distance of a prong from the edge of the detector as the minimal distance from all the prong’s 3D points to a detector edge. It is required that either the distance of both prongs is > 5 cm from the edge or else that the combined distance of both prongs is > 15 cm from the edge.

In order to associate a prong with the SparseSSnet identified pixels, the 3D prong is projected onto the 2D images described above. The pixels in this projection are then matched to the prong.

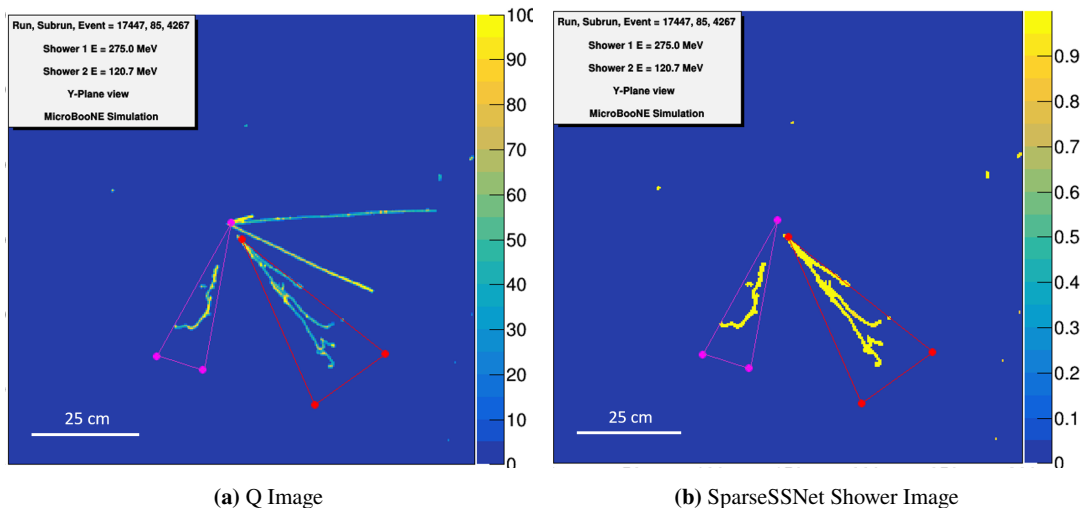


Figure 1. Event displays of a simulated CC π^0 event. (a) shows the raw Q image and (b) shows the SparseSSNet shower score with track-like particles masked out. In both (a) and (b) the leading reconstructed photon is represented by the red triangle and the sub-leading reconstructed photon (shower 2) is represented by the magenta triangle.

The kinetic energies of track-like prongs are calculated using the track lengths. The direction of track-like particles is also reconstructed as described in ref. [4]. Energy deposits from electrons and photons will suffer gaps due to radiated photons which the 3D track reconstruction cannot handle. Shower-like prongs therefore need a separate reconstruction as described in section 3.

3 Electromagnetic shower energy reconstruction

The electromagnetic shower reconstruction algorithm is run to find any associated shower particles and reconstruct their kinetic energies once a candidate vertex is isolated. The method described here builds on a previous MicroBooNE shower reconstruction described in ref. [13]. The algorithm described retains many vital features of the previous version, especially the use of a semantic segmentation neural network for pixel labeling. Important updates have been added including many simplifications of the algorithm made possible by the improved SparseSSnet.

The first step of the reconstruction is to mask the image to only use the pixels identified as shower by SparseSSNet. SparseSSNet outputs a value for each pixel indicating how likely it is that the pixel is part of a shower. This value, called shower score, falls between 0.0 and 1.0 where 1.0 indicates a shower-like pixel. Pixels with a shower score of > 0.5 and intensity $> 10 Q$ (charge) counts are kept, all other pixels are masked out for the rest of the shower reconstruction. A threshold of $10 Q$ is used to remove wire noise. This cut is standard among all tools used in this analysis. The value was chosen based on the distribution of Q from minimum ionizing particles (MIPs), which peaks at $\approx 40 Q$. A template isosceles triangle is then placed with its apex at the reconstructed vertex position, pointing in the positive wire direction. The triangle is optimized to choose the shower direction, length, and opening angle for which the triangle contains the most pixels with non-zero charge. These parameters each start at the minimum value shown in table 1. In order to allow for

showers that are detached from the vertex, a gap parameter is introduced allowing the triangle to start further from the vertex.

Once a first shower candidate has been found by the reconstruction algorithm, the pixels found in the shower are masked out. If the total amount of charge remaining passing the cuts of a shower score of > 0.5 and intensity $> 10 Q(\text{charge})$ is $> 5000 Q$, the second shower algorithm is run on the masked image. The total range of allowed values for each of the template triangle parameters is shown in table 1. In this table “first shower” refers to the shower found in the first pass of the shower reconstruction which is aimed at finding showers near the reconstructed vertex. “Second shower” refers to the shower found on the second run of the algorithm and has expanded parameters to search for detached showers as described in the following paragraphs. The parameters in each case are optimized sequentially in the order of direction, gap size, opening angle, and length.

Table 1. Range of parameters for the shower reconstruction algorithm. Parameters are changed in the second shower search to allow for the capture of showers detached from the vertex.

	Minimum value	Maximum value
Direction	0 degrees	360 degrees
Opening angle	17 degrees	75 degrees
Length (first shower)	3 cm	35 cm
Length (second shower)	3 cm	60 cm
Gap Size (first shower)	0 cm	17 cm
Gap Size (second shower)	0 cm	90 cm

Figure 1 shows an example display demonstrating the 2D shower reconstruction on a simulated CC π^0 event. The SparseSSNet shower-like particles are shown in (b) with the track-like particles masked out. The final optimized showers are shown in red (first shower) and magenta (second shower). In this example, the algorithm found the proper gap of the first shower, but not the second shower. This is acceptable as gap size is not a value that is utilized in any other part of the DL analysis. The reconstructed energy of each shower is also reported. Energy reconstruction of electromagnetic showers is a crucial component of the LEE analysis this work supports, which is designed to measure electrons from CC neutrino interactions over a broad energy range from 35 to 1200 MeV. This analysis aims to develop and validate a energy reconstruction procedure for showers in this broad energy range. To determine the energy of each shower, the charge of all shower pixels enclosed in the Y -view triangle is integrated. This total shower charge will hereby be denoted by Q_{sh} . We use the Y -view, which is the collection plane, because it has the highest signal-to-noise ratio of the three planes [9].

To convert the reconstructed charge to energy, the Q_{sh} in a sample of simulated events is compared to the generated energy of the electrons in the events. A Q_{sh} -to-MeV conversion line is determined as described below and shown in figure 2. As the focus of the larger analysis is $1e1p$ events, a sample of simulated $1e1p$ events are used. The simulated electron energy is plotted versus the reconstructed Y -view total shower charge sum for events in a ν_e simulation sample selected by the MicroBooNE DL-based $1e1p$ analysis. In each vertical bin, the peak is found with the help of a Gaussian distribution (represented by the black points in figure 2). The edge bins which have smaller statistics are excluded. Two examples of the Gaussian fits are shown in figure 3. While the

Gaussian fits capture the bulk of simulated events well, one can see a tail in each distribution out to higher simulated electron energies which is not captured by the Gaussian. This is expected, as these tails correspond to electron showers which are not fully reconstructed (i.e., showers which pass through unresponsive wires or exit the active volume). The points are then fit to a line. The slope of this line is used for the Q_{sh} -to-MeV conversion and is referred to in the rest of this article as m_{e^-} . While a Gaussian is not a perfect fit to the data in, the use of m_{e^-} , derived from these points, will be validated in section 4 and section 5. The resulting equation is:

$$\text{Electron: } E \text{ [MeV]} = (1.26 \pm 0.01 \times 10^{-2}) \times Q_{\text{sh}} \text{ [Q Counts]}. \quad (3.1)$$

where the error corresponds to the uncertainty on the linear fit, which represents the statistical error on the simulated electron sample used for the fit.




Figure 2. Simulated electron energy vs Q_{sh} for a sample of generated $1e1p$ events. The linear fit is used in the shower energy calculation.

Various detector effects could cause this fit value to change. Specifically the amount of detected energy will depend on the position, amount of energy deposited, and the orientation of the particle’s trajectory with respect to the wires [10, 11, 14]. Additionally, the value of m_{e^-} derived here assumes good argon purity and could be affected by periods of low purity. Potential systematic uncertainty could be introduced by these effects. The results shown in section 6 give an estimate of the size of the detector effect.

Using this shower energy calculation, we look at the energy resolution for a sample of simulated CC ν_e events containing an electron and no final state π^0 . This isolates electrons from photons which will be discussed further in section 4.1. The following selection criteria are used:




Figure 3. Example distributions of simulated electron energies (solid lines) and corresponding Gaussian distributions (dashed lines) within two different shower charge sum ranges. The peaks of the Gaussian fits are used to generate the black points in figure 2.

1. Reconstructed vertex is less than 5 cm from simulated neutrino interaction vertex;
2. One simulated electron contained in event;
3. No final state π^0 ;
4. One reconstructed shower; and
5. 1e1p Boosted Decision Tree (BDT) score is greater than 0.7 [3].

The energy resolution, defined here as:

$$E_{\text{res}} = \frac{E_{\text{reco}} - E_{\text{sim}}}{E_{\text{sim}}} \quad (3.2)$$

for this sample of simulated events is seen in figure 4. The mean is at -0.07 and the RMS is 0.22.

The shower energy reconstruction presented here will be utilized and validated in section 4 and section 5. As stated earlier, the shower energy is calculated using a MicroBooNE simulation sample comprising of simulated neutrino events overlay with off-beam cosmic ray data. It is therefore notable that the energy calculated with this sample works well when applied to data samples as




Figure 4. The energy resolution for a sample of simulated electrons as described by eq. (3.2). The y axis has the raw number of simulated events without scaling. The dashed vertical line is included at $E_{\text{res}} = 0.0$ for reference.

shown in the next sections. It will be seen that the π^0 and Michel e^- samples both have good data/simulation agreement using eq. (3.1). A further test is done for each sample that finds the charge-to-energy conversion factor that gives the best agreement to known physical values (the π^0 rest mass and the Michel electron spectrum cutoff). This test shows great data/simulation agreement with each sample. The results are also comparable to m_{e^-} given in eq. (3.1), even at different energy scales, validating the use of this linear conversion factor.

4 Neutral pion sample

The first sample used to analyze the performance of the shower reconstruction is a sample of π^0 events. The identification and reconstruction of π^0 events is presented. This is followed by both a verification of the Q_{sh} -to-MeV conversion value (m_{e^-}) agreement to data and simulation, and a verification of agreement between data and simulation which is accomplished by using a well-measured physical quantity: the π^0 invariant mass (135 MeV).

4.1 Identification and reconstruction

The reconstruction of π^0 events relies on the shower reconstruction described in section 3 and introduces 3D shower reconstruction. The starting point is either a track-track or track-shower

vertex. The shower reconstruction is then applied as discussed above, centering on clusters of electromagnetic charge, but not requiring that the charge be attached to the starting vertex. This leads to the bulk of π^0 events selected for this analysis matching two specific topologies. The most prevalent π^0 topology in this study is from charge current (CC) π^0 where the scattered muon and the proton from the Δ decay form the vertex, and there are two disconnected electromagnetic showers from π^0 decays. The second is from neutral current (NC) π^0 where one photon converted within the 0.3 cm wire spacing and thus forms a vertex with the proton, while the second photon is displaced from the vertex. As a result, contributions to the π^0 selection discussed here will come from NC π^0 and CC π^0 , as well as CC π^- and CC π^+ events where the π^+/π^- undergoes charge exchange within the nucleus.

To fully calculate the kinematics of the π^0 event we must determine both the energy and the 3D angle of the electromagnetic showers. This is essential for reconstructing the π^0 invariant mass — a useful quantity to test the shower reconstruction described in section 3. In order to cluster a full 3D shower, the 2D projections on the different wire planes are compared for overlap in time. The overlap fraction is defined as the fraction of shower pixels in the collection plane shower that overlap in time with shower pixels from a 2D shower in another plane, in which the U and V planes are considered separately. If the overlap fraction is > 0.5 in either or both planes, the pixels that overlap between the collection plane shower and the shower in another plane with the highest overlap fraction are used to calculate a cluster of 3D shower points. The direction is found by using the calculated center of the 3D point cluster and the event vertex.

This 3D reconstruction leads to what is referred to as the “ π^0 pre-selection cuts”. These are requirements that are necessary in order to reconstruct two 3D showers.

1. Vertex passes fiducial volume containment requirement (described in section 2);
2. Two collection plane showers, each with reconstructed energy greater than 35 MeV;
3. Both collection plane showers have an overlap fraction with a shower in another plane greater than 0.5;
4. If a collection plane shower matches with showers in both the U and V planes, the one with the highest overlap fraction is chosen; and
5. The two collection plane showers cannot match to the same shower in another plane.

After the “ π^0 pre-selection cuts” have been applied, a substantial number of selected events have showers that are mis-reconstructed. The sample also contains a large number of backgrounds such as cosmic muons. To improve the selection, the following requirements are introduced that reduce mis-reconstruction and remove backgrounds. These are referred to as “box cuts” as they are hard cuts designed to remove background events at the tails of the distributions of various variables.

One of the variables used in these “box cuts” is a Δ mass test variable. This value is calculated for all those events passing the “ π^0 pre-selection cuts”. The 4-vector of the reconstructed showers are used along with the 4-vector of the proton-like prong. The reconstruction of the proton-like prong is discussed in ref. [4]. These three objects are assumed to have come from a Δ decay and are therefore used to reconstruct a Δ rest mass. The tails of this distribution are comprised of mis-reconstructed π^0 events and cosmic muon backgrounds which allows for another box cut. The box cuts are then:

1. Reconstructed π^0 mass is less than 400 MeV;
2. Reconstructed energy of the leading photon is greater than 80 MeV;
3. The charge sum of all pixels (both track and shower) within 2 cm of the vertex is greater than 250 Q counts;
4. Leading shower reconstructed angle w.r.t. beam direction is less than 1.5 radians;
5. The angle between the two photons is less than 2.5 radians; and
6. Δ mass test variable is between 1000 and 1400 MeV.

Here, “leading photon” refers to the simulated photon with the highest energy. The reconstructed leading shower is the shower with the highest Q_{sh} . An additional requirement on the 1e1p Boosted Decision Tree (BDT) of < 0.7 is further added at this stage to maintain blindness to the LEE for this study as required by the MicroBooNE blindness procedure [3]. This BDT is used to select events with one electron and one proton, so this is a relatively small cut that is a flat $\approx 3\%$ effect. Events above this threshold may be mis-reconstructed 1e1p events to which we are currently maintaining blindness.

Using the shower energy calculation from the electron sample on both electron and photon showers assumes that the energy of both shower types is reconstructed the same way. To demonstrate that this is valid, the same simulated energy vs Q_{sh} plot and fit is performed on two samples of photons: leading and sub-leading photons from a sample of simulated $\text{CC}\pi^0$ events. To ensure this fit is performed only over well-reconstructed events, the π^0 selection box cuts are applied. The results are shown in figure 5. The resulting fit equations are:

$$\text{Leading Photon: } E \text{ [MeV]} = (1.25 \pm 0.02 \times 10^{-2}) \times Q_{\text{sh}} \text{ [Q Counts]}. \quad (4.1)$$

$$\text{Sub-leading Photon: } E \text{ [MeV]} = (1.20 \pm 0.02 \times 10^{-2}) \times Q_{\text{sh}} \text{ [Q Counts]}. \quad (4.2)$$

It is seen from the fit results that the leading photon fit very closely matches the electron fit as expected while the sub-leading photon fit does not match as closely. This is due to the worse reconstruction in the sub-leading photon. Therefore, the reconstruction uses the charge-to-energy (Q_{sh} -to-MeV) conversion value found for the simulation electron sample, m_{e^-} .

For the purpose of studying this sample, Monte Carlo simulation has been broken into various categories. “NC π^0 ” are neutral current π^0 events with a well reconstructed vertex, which is a vertex within 5 cm of the true generated vertex. “CC π^0 ” are defined similarly for charged current π^0 . “Offvtx π^0 ” are π^0 events with poorly reconstructed vertices. Here, off-vertex means that the reconstructed vertex is further than 5 cm from a true generated neutrino vertex. “Non π^0 ” events are broken into on and off vertex as well. ν_e events are all events that originated from a ν_e . Cosmic background events also remain after selection.

Simulated energy resolution of the photons is presented in figure 6, where resolution is defined in eq. (3.2). This plot is made using the specialized high POT simulation sample of containing only events with a π^0 in the final state. For the purpose of scaling the different samples and background contributions, the events are POT scaled to match the total data exposure. This distributions in this plot have all of the π^0 selection cuts applied. The energy resolution is shown separately for leading (highest energy) photon and sub-leading photon.

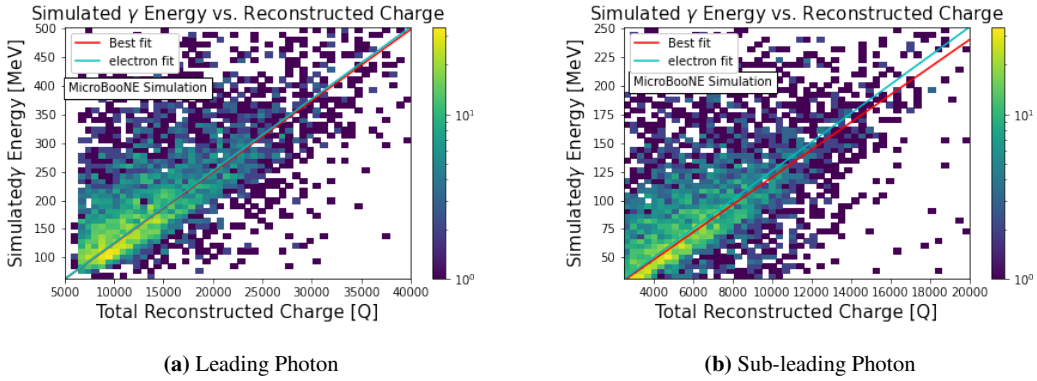


Figure 5. Simulated photon energy vs Q_{sh} for a sample of generated CC π^0 events with the π^0 selection applied. The best fit is shown for this sample as well as the best fit from the electron fit. (a): leading photon, (b): sub-leading photon.




Figure 6. The energy resolution for each of the decay photon in the selected π^0 sample. The leading photon is shown in (a) and the sub-leading photon is shown in (b). The events have been scaled to match the total data POT of 6.67×10^{20} . Resolution is defined in eq. (3.2). The dashed vertical line is included at $E_{res} = 0.0$ for reference.

Table 2 shows the mean and RMS of the distributions shown in figure 6. Both the leading and sub-leading photon E_{res} mean is close to zero. The leading photon has a smaller RMS than the sub-leading photon, which is more broad and has a larger tail. The resolution of photons from π^0 is worse than that seen in the electrons, but they have similar bias as indicated by the mean. There are two main causes for the difference between leading and sub-leading photons. The first is that the sub-leading photon is generally the lower energy of the two. Failing to reconstruct a small number of pixels will have a larger effect on a lower energy shower which has fewer true shower pixels associated with it. The other cause are events where the leading and sub-leading shower are close together. Part of the sub-leading shower is reconstructed as part of the other shower as the leading shower is prioritized. Of these two sub-leading photon reconstruction failure modes, the first is dominant. Some of the other mis-reconstructed photons in both the leading and sub-leading plot are caused by mistakes in SparseSSNet, overlapping cosmic rays that were not removed properly, and unresponsive wires in the collection plane.

Table 2. Characterization of the resolution distributions shown in figure 4, figure 6, and figure 8.

	Mean E_{sim}	Mean E_{res}	RMS of E_{res}
Electron	744.8 MeV	-0.07	0.22
Leading Photon	230.3 MeV	-0.07	0.36
Sub-leading Photon	98.8 MeV	0.04	0.69
	Mean θ_{sim}	Mean θ_{res}	RMS of θ_{res}
Between two photons	54.1°	0.05	0.50

Figure 7 shows a data to MC simulation comparison of the reconstructed photon energies in the π^0 sample. The uncertainty bars on the simulation distribution represent the statistical uncertainty. The total number of simulation events has been scaled to match the total number of data events. This plot uses both the MicroBooNE overlay simulation samples combined with the high POT π^0 simulation sample and data as described in section 3. The spectra seen in this figure are at higher reconstructed energies than the Michel e^- in section 5 and closely mirror the shower energy scale of the MiniBooNE LEE. The χ^2 which is reported in the caption of this plot is a combined Neumann-Pearson (CNP) χ^2 [15]. The χ_{CNP}^2 is defined as:

$$\chi_{\text{CNP}}^2 = \sum_i \begin{cases} \frac{(\mu_i - M_i)^2}{\frac{3}{1/M_i + 2/\mu_i}} & M_i \neq 0 \\ \frac{(\mu_i - M_i)^2}{\frac{\mu_i}{2}} & M_i = 0 \end{cases} \quad (4.3)$$

where μ_i and M_i are the number of predicted and observed events in a given bin.

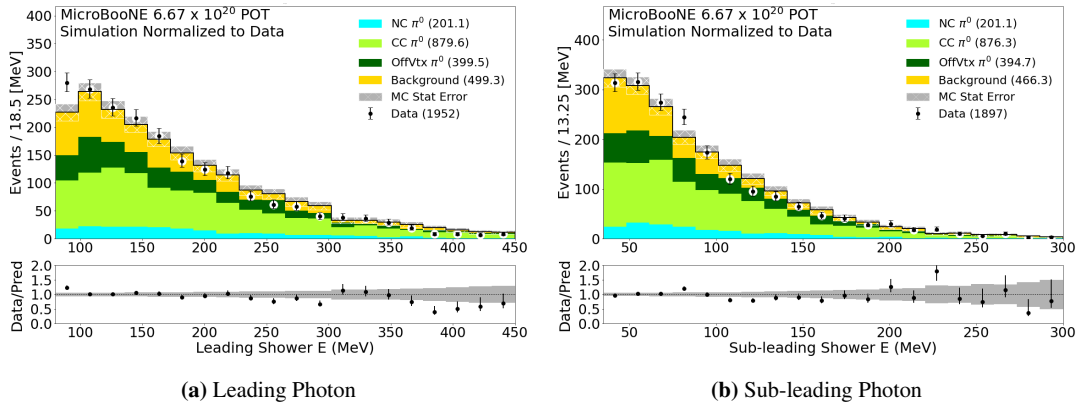


Figure 7. The reconstructed photon energies for events passing all selection cuts. The leading photon is shown in (a) and the sub-leading photon is shown in (b). The MC simulation samples have been normalized to the total number of data events. The data events are shown by black points. The number of events in each category is shown in the legend in parentheses. The $\chi_{\text{CNP}}^2/19(\text{dof}) = 1.267$ with a p-value of 0.193 for the leading shower and the $\chi_{\text{CNP}}^2/19(\text{dof}) = 0.973$ with a p-value of 0.491 for the sub-leading shower.

Figure 8 shows the resolution of θ . θ is the opening angle between the two showers, either simulated or reconstructed, used in the π^0 mass reconstruction. The opening angle resolution is defined as:

$$\theta_{\text{res}} = \frac{\theta_{\text{reco}} - \theta_{\text{sim}}}{\theta_{\text{sim}}}. \quad (4.4)$$

This figure uses the same simulation sample with the same selection requirements as figure 6. The distribution is characterized in table 2. The mean is close to zero indicating little bias.




Figure 8. The opening angle (θ) resolution of the decay photons in the selected π^0 sample. The events have been scaled to match total data POT of 6.67×10^{20} . Resolution is defined in eq. (4.4). The dashed vertical line is included at $\theta_{\text{res}} = 0.0$ for reference.

The π^0 rest mass can now be reconstructed using the following equation:

$$M_{\pi}^0 = \sqrt{4 \sin^2 \left(\frac{\theta}{2} \right) (E_1)(E_2)} \quad (4.5)$$

where E_1 is the leading photon energy and E_2 is the sub-leading photon energy. The result of this reconstruction is shown in figure 9. As in figure 7, the total number of simulation events has been scaled to match the total number of data events. The distributions of both data and simulation peak around 135 MeV, which is the accepted π^0 rest mass.

4.2 Validation of agreement between simulation, data, and true rest mass

The π^0 sample is next used to verify the agreement in data and simulation of the shower energy scale using the known π^0 invariant mass M_{π^0} . Test points are found representing the best-fit of




Figure 9. The calculated π^0 mass for events passing all selection cuts. The MC simulation samples have been normalized to total number of data events. The data events are shown by black points. The number of events in each category is shown in the legend in parentheses. The $\chi^2_{\text{CNP}}/19(\text{dof}) = 0.976$ with a p-value of 0.486 for the MC prediction.

the Q_{sh} -to-MeV conversion factor (m) to $M_{\pi^0} = 135$ [MeV/ c^2]. This is done for a sample of each simulation and data. The goal for each is to find the value of m that yields a π^0 mass distribution that peaks closest to the true value of 135 [MeV/ c^2]. This value of m is then compared between data and simulation and to the electron m_{e^-} value found in section 3.

To find the optimal m , the following χ^2 formula is minimized:

$$\chi^2 = \sum_i \left(\frac{(135[\text{MeV}/c^2] - M_{\pi^0}^i)}{dM} \right)^2. \quad (4.6)$$

where i is each π^0 event in the given sample, dM is 29.8 [MeV/ c^2] based on the width of a Gaussian fit to the good simulation π^0 distribution (the NC π^0 and CC π^0 categories in figure 9). $M_{\pi^0}^i$ represents the π^0 mass and is given in this case by:

$$M_{\pi^0}^i = \sqrt{4 \sin^2 \left(\frac{\theta}{2} \right) (m \times (Q_{\text{sh}})_1)(m \times (Q_{\text{sh}})_2)} \quad (4.7)$$

where Q_{sh} is the reconstructed shower charge and θ is the reconstructed opening angle between the two showers. The same χ^2 formula is minimized for both MC simulation and data. The 1- σ range is

calculated on these fit points by looking at the range of m values for each data and simulation which give a χ^2 value satisfying the Wilks' theorem condition $|\chi^2(m) - \min_m[\chi^2(m)]| < 1$ [16]. The χ^2 distributions can be found in figure 10. The resulting values can be seen in the first two rows of table 3. Excellent agreement is seen between the data and the simulation best fit points indicating that the simulation derived shower conversion factor is valid to use on data.

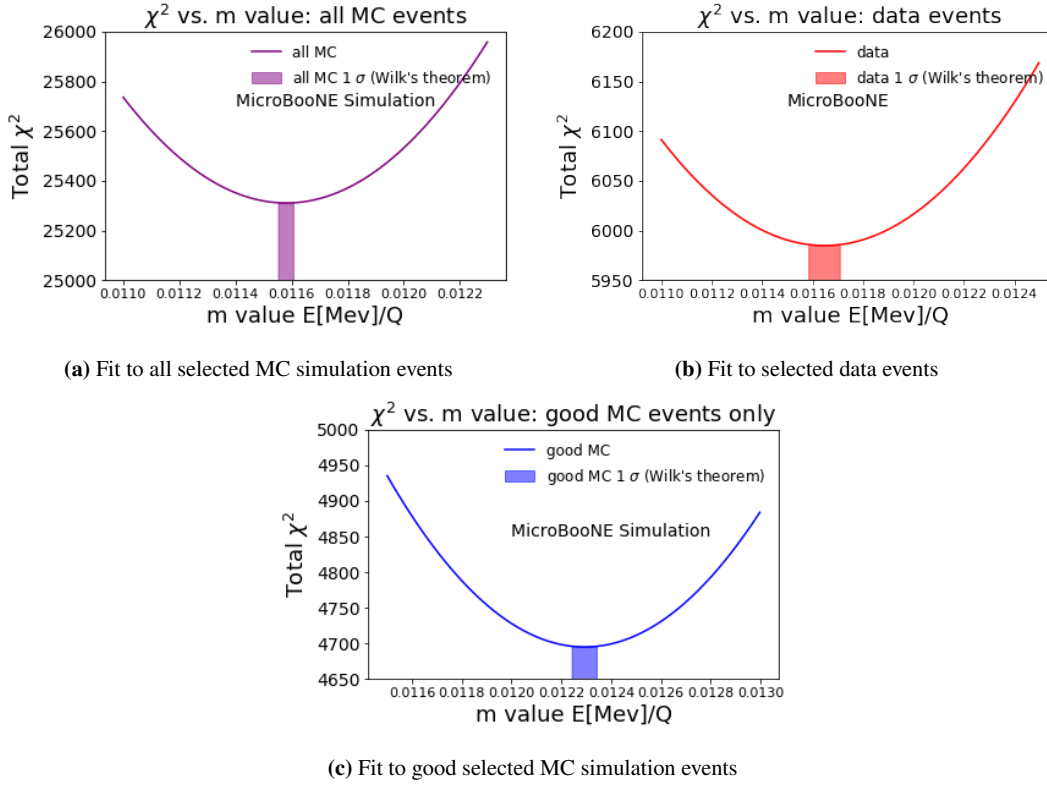


Figure 10. Total χ^2 vs. m distributions for all MC simulation (a), data (b), and good MC simulation (c) that pass the π^0 selection criteria.

Another important consideration is how closely the best fit m values derived here match m_e in eq. (3.1). The best fit value of m derived from the π^0 sample has the potential to be affected by many factors. The largest of these factors is the amount of background events selected. The π^0 selection presented in section 4.1 contains many background events, which should not necessarily reconstruct as a π^0 mass of 135 [MeV/c²]. To account for this background in simulation a second fit is performed only using good simulation events. In this instance good simulation is defined as simulated events that pass all π^0 cuts, have a simulated final state π^0 , and have a reconstructed vertex within 5 cm of the true simulated neutrino interaction vertex. These are the only types of events in the selection which should result in a reconstructed π^0 mass of 135 MeV. A further cut of π^0 mass < 200 MeV is added to prevent mis-reconstructed events from having a large effect on the χ^2 .

To account for the background events in data, the optimal m found previously is shifted by the same amount that the MC m is shifted when backgrounds are included (7.11×10^{-4}) as seen

the last two rows in table 3. Data (un-shifted) are the results from minimizing eq. (4.6) over all data points. Data (shifted) shows the data points applying the shift found by comparing the fit over all simulation to the fit over good simulation. This shift retains the excellent data and simulation agreement. The result on data only differs by 1.6% from the value seen in eq. (3.1). This indicates that, at the photon energy scale seen in the π^0 sample, the charge-to-energy conversion factor is valid. Data and simulation agreement of these results and agreement to eq. (3.1) is discussed further in section 6 and examined in combination with the results from the Michel e^- sample.

Table 3. The best value of m (MeV/ Q_{sh}) for each data and MC simulation sample and the range found using Wilks’s theorem. Results are shown before accounting for background (top two rows) and after (bottom two rows).

Sample	m [MeV/Q]	m range	χ^2/NDF
All MC	1.159×10^{-2}	$[1.156 \times 10^{-2}, 1.160 \times 10^{-2}]$	$25310.3/9473 = 2.7$
Data (un-shifted)	1.165×10^{-2}	$[1.159 \times 10^{-2}, 1.170 \times 10^{-2}]$	$5984.7/1973 = 3.0$
Good MC	1.230×10^{-2}	$[1.225 \times 10^{-2}, 1.234 \times 10^{-2}]$	$4694.9/3039 = 1.5$
Data (shifted)	1.236×10^{-2}	$[1.230 \times 10^{-2}, 1.241 \times 10^{-2}]$	$5984.7/1973 = 3.0$

5 Michel electron sample

The second sample used to validate the shower energy reconstruction consists of Michel electrons from decays of ν_μ -sourced stopped muons in the detector. This is the first time a Michel sample has been reconstructed in a LArTPC which is dominated by Michels from ν_μ interactions. As in section 4, we begin by describing the event selection criteria for this sample. Next, we examine the data/simulation agreement in the Michel shower energy spectrum. Finally, we assess the agreement of the Michel sample with the physical Michel cutoff of $m_\mu/2 = 52.8$ MeV through a fit procedure. The data and MC simulation results of this fit agree, validating the use of the ν_e simulation-derived Q_{sh} -to-MeV on data. Both the data and ν_e simulation fits also show consistency between the simulation-derived Q_{sh} -to-MeV conversion value validating the absolute shower energy scale of the DL-based analysis in the low energy region ($\lesssim 50$ MeV).

5.1 Identification and reconstruction

The Michel electron sample has been chosen in order to validate the reconstructed energy scale of lower-energy electrons, slightly below the energy scale of electrons in the low-energy excess search. Previous work has been performed in MicroBooNE using a larger sample than presented here, as seen in ref. [17]. The study presented here uses a different selection and is designed to test the reconstruction of showers used in the DL low-energy excess search. However, as shown below, our results are consistent with those shown in [17] in both the reconstructed Michel energy spectrum and corresponding energy resolution.

Muon-Michel vertices are identified through the following requirements:

1. Two prongs at the vertex;
2. Long prong track-length > 100 cm (candidate muon);

3. Short prong track-length < 30 cm (candidate Michel);
4. Long track consists of < 20% SparseSSNet shower-like pixels (candidate muon);
5. Short track consists of > 80% SparseSSNet shower-like pixels (candidate Michel); and
6. $\phi_\mu < 0.5$ radians.

where ϕ_μ is the azimuthal angle of the muon with respect to the horizontal plane, where $\phi_\mu = 0.5$ rad corresponds to downward-going muons. Lastly, in events with more than one selected vertex, we keep only the first vertex. This strategy was chosen as it does not bias the Michel energy spectrum.




Figure 11. ϕ_μ distribution for selected events in both data and MC simulation, corresponding to $\approx 5.3 \times 10^{19}$ POT. The selection cut requiring $\phi_\mu < 0.5$ radians is indicated by the dotted line. The MC simulation samples have been normalized to total number of data events. The data events are shown by black points. The number of events in each category is shown in the legend in parentheses. The uncertainty bars here are statistical only. The $\chi^2_{\text{CNP}}/9(\text{dof}) = 0.822$ with a p-value of 0.596 for the MC prediction.

There are two types of Michel electrons that are isolated before the selection. The first are Michels in neutrino events, which can be compared between simulation and data. The second are Michel electrons from stopped cosmic muons. As our simulation samples contain simulated neutrino events overlaid with cosmic data, all of the Michels from stopped cosmic muons come from actual data. Therefore, the only Michels that are truly simulated are those on simulated neutrino events. It is important for these studies to be certain that the Michels in the simulation sample do not come from the stopped cosmic muons. This ensures we are comparing Michels in neutrino events from data to simulated Michels. This is achieved by the final requirement on the muon polar angle ϕ_μ , which removes a majority of the predominately downward-going cosmic muons as shown in

figure 11. As discussed in section 2, the DL vertices search for the intersection of two “prongs”. A muon decay into a Michel electron forms this pattern and is therefore often reconstructed at the initial vertex stage of the event reconstruction.

For Michel electrons in events passing these cuts, the electromagnetic shower reconstruction algorithm from section 3 is applied to find Q_{sh} . This is then converted to a shower energy via eq. (3.1). Figure 12 shows the shower energy distribution for Michel electrons in both data and simulation. Note that data here come from an open beam data set of corresponding to $\approx 5.3 \times 10^{19}$ POT. The simulated events are broken into various categories to indicate which type of event caused the muon. The majority of events come from ν_μ interactions within the active detector volume. Some events from cosmic muons and muons from ν_μ interactions outside the active volume remain in the sample.




Figure 12. Electron energy distribution for Michels in both data and MC simulation after all selection criteria have been applied, corresponding to $\approx 5.3 \times 10^{19}$ POT. The MC simulation samples have been normalized to total number of data events. The data events are shown by black points. The number of events in each category is shown in the legend in parentheses. The uncertainty bars here are statistical only. The $\chi^2_{\text{CNP}}/9(\text{dof}) = 0.608$ with a p-value of 0.857 for the MC prediction.

One can see that the high-end tails for the shower energy distribution in both data and simulation fall off around 60 MeV as expected. While not as sharp as the cut-off in ref. [17], the results are consistent. The high energy tail above the true value of 52.8 MeV likely is due to over estimation of shower energy reconstruction seen in figure 4. The shower energies of this sample are much lower than those seen in the π^0 sample. The good data/simulation agreement within statistical uncertainty indicates that the shower algorithm performs well down to low energies. This agreement will be quantified further in the next sections.

5.2 Validation of agreement between simulation, data, and true Michel cutoff

We now perform a validation of the absolute shower energy scale and analyze the data/simulation agreement with the Michel sample analogous to the π^0 study described in section 4.2. In the case of the Michel cross-check, we fit the Q_{sh} (reconstructed shower charge) spectrum shown in figure 13 to the following five-parameter function:

$$f\left(x = \frac{Q_{\text{sh}}}{Q_{\text{cutoff}}}; \sigma, N\right) = N \int_0^1 (3y^2 - 2y^3) \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-y)^2}{2\sigma^2}\right) dy \quad (5.1)$$

where:

$$\sigma = y \sqrt{r_1^2 + \frac{r_2^2}{yQ_{\text{cutoff}}} + \left(\frac{r_3}{yQ_{\text{cutoff}}}\right)^2}. \quad (5.2)$$

Here, $f(x)$ represents a parameterization of the true Michel spectrum convoluted with a Gaussian representing charge resolution [18]. N is a floating normalization parameter, and $\{r_1, r_2, r_3\}$ represent contributions to the charge resolution corresponding to a constant noise term, a statistical charge-counting term, and a Gaussian noise term, respectively. Q_{cutoff} represents the cutoff of the Michel shower energy spectrum, which, after the Q_{sh} -to-MeV conversion, should correspond to the Michel energy cutoff of $m_\mu/2 \approx 52.8$ MeV. The integration over the variable y represents a scan over the simulated shower charge spectrum. The expression is invariant under $\int_0^1 I(y) dy \rightarrow Q_{\text{cutoff}}^{-1} \int_0^{Q_{\text{cutoff}}} I(Q_{\text{sh}}^*/Q_{\text{cutoff}}) dQ_{\text{sh}}^*$, where $I(\dots)$ represents the integrand in eq. (5.1).

We first fit $f(x)$ to the Michel spectrum in data and simulation by varying all five parameters $Q_{\text{cutoff}}, N, r_1, r_2, r_3$. This is done by minimizing the χ^2 :

$$\chi^2(Q_{\text{cutoff}}, N, r_1, r_2, r_3) = \sum_i \left(\frac{(O_i - f(x = \frac{Q_{\text{sh}})_i}{Q_{\text{cutoff}}}; N, r_1, r_2, r_3)}{\sigma_{i,\text{stat.}}} \right)^2 \quad (5.3)$$

where O_i is the number of observed Michel events in Q_{sh} bin i and $\sigma_{i,\text{stat.}} = \sqrt{O_i}$ is the Poisson error. There are 12 shower charge bins ranging from 0–6000 counts in this fit. Figure 14 shows 2D confidence regions for Q_{cutoff} versus the different resolution parameters. They are calculated by fixing the remaining three parameters at their best fit values and using Wilks's theorem for two free parameters. As shown in figure 14, the fit generally prefers a large contribution from the flat resolution term r_1 . In data, one can see that the 1σ r_1 contour prefers a fractional resolution of ≈ 0.3 , while the 1σ r_2 and r_3 contours are both consistent with zero. In simulation, the 1σ r_1 contour prefers a fractional resolution of ≈ 0.25 . The 1σ r_3 contour is consistent with zero, but the 1σ r_2 contour prefers a value of ≈ 13 [Q counts] $^{1/2}$. This turns out to be a similar contribution when compared to the flat r_1 term. A near-flat energy resolution for Michel showers is consistent with previous MicroBooNE work [17].

Next, the parameters $\{N, r_1, r_2, r_3\}$ are fixed at the minimum χ^2 values (respectively for data/simulation) and a χ^2 scan over Q_{cutoff} is performed, this time only including the tail of the Q_{sh} spectrum ($Q_{\text{sh}} > 3000$ Q counts). The purpose of this one-dimensional scan is to obtain the χ^2 minimum and corresponding 1σ interval on Q_{cutoff} using Wilks's theorem for one fit parameter [16]. Figure 13 shows the observed Michel shower charge spectra in data and simulation along with their respective best fits from the 1D scan. Figure 15 shows χ^2 as a function of Q_{cutoff} . The Wilks's




Figure 13. Top: Michel shower charge sum spectrum in data and MC simulation along with the corresponding best fit to eq. (5.1) (allowing only Q_{cutoff} to vary in the fit). This sample corresponds to $\approx 5.3 \times 10^{19}$ POT. **Bottom** Ratio of the data/simulation to the corresponding fit. The MC simulation and fit result here have been normalized to match the data.

theorem 1σ interval on Q_{cutoff} corresponds to the points for which $\chi^2(Q_{\text{cutoff}}) - \min\{\chi^2; Q_{\text{cutoff}}\} \leq 1$. In order to get the charge to energy conversion factor m from Q_{cutoff} , $m = \frac{52.8 \text{ MeV}}{Q_{\text{cutoff}}}$. The best fit and 1σ intervals for m are given in table 4 along with the χ^2/NDF of the best fit. One can see excellent agreement between data and simulation, demonstrating the consistency of the shower reconstruction. Note that the 1σ interval on Q_{cutoff} is larger in data than in simulation — this is because the statistical error on the data Michel sample is larger than that on the simulated Michel sample. The data/simulation agreement and consistency with eq. (3.1) demonstrated by this study are discussed further in section 6 in combination with the results from the π^0 sample.

Table 4. The best fit values and 1σ ranges (via Wilks’ theorem) for m along the χ^2/NDF of that fit given by eq. (5.3) for both data and MC simulation. The fit here is the one-dimensional scan over Q_{cutoff} transformed into m as described in the text.

Sample	m [MeV/Q]	m range	χ^2/NDF
Data	1.341×10^{-2}	$[1.282 \times 10^{-2}, 1.401 \times 10^{-2}]$	$2.17/6 = 0.4$
MC	1.308×10^{-2}	$[1.279 \times 10^{-2}, 1.334 \times 10^{-2}]$	$2.73/6 = 0.5$

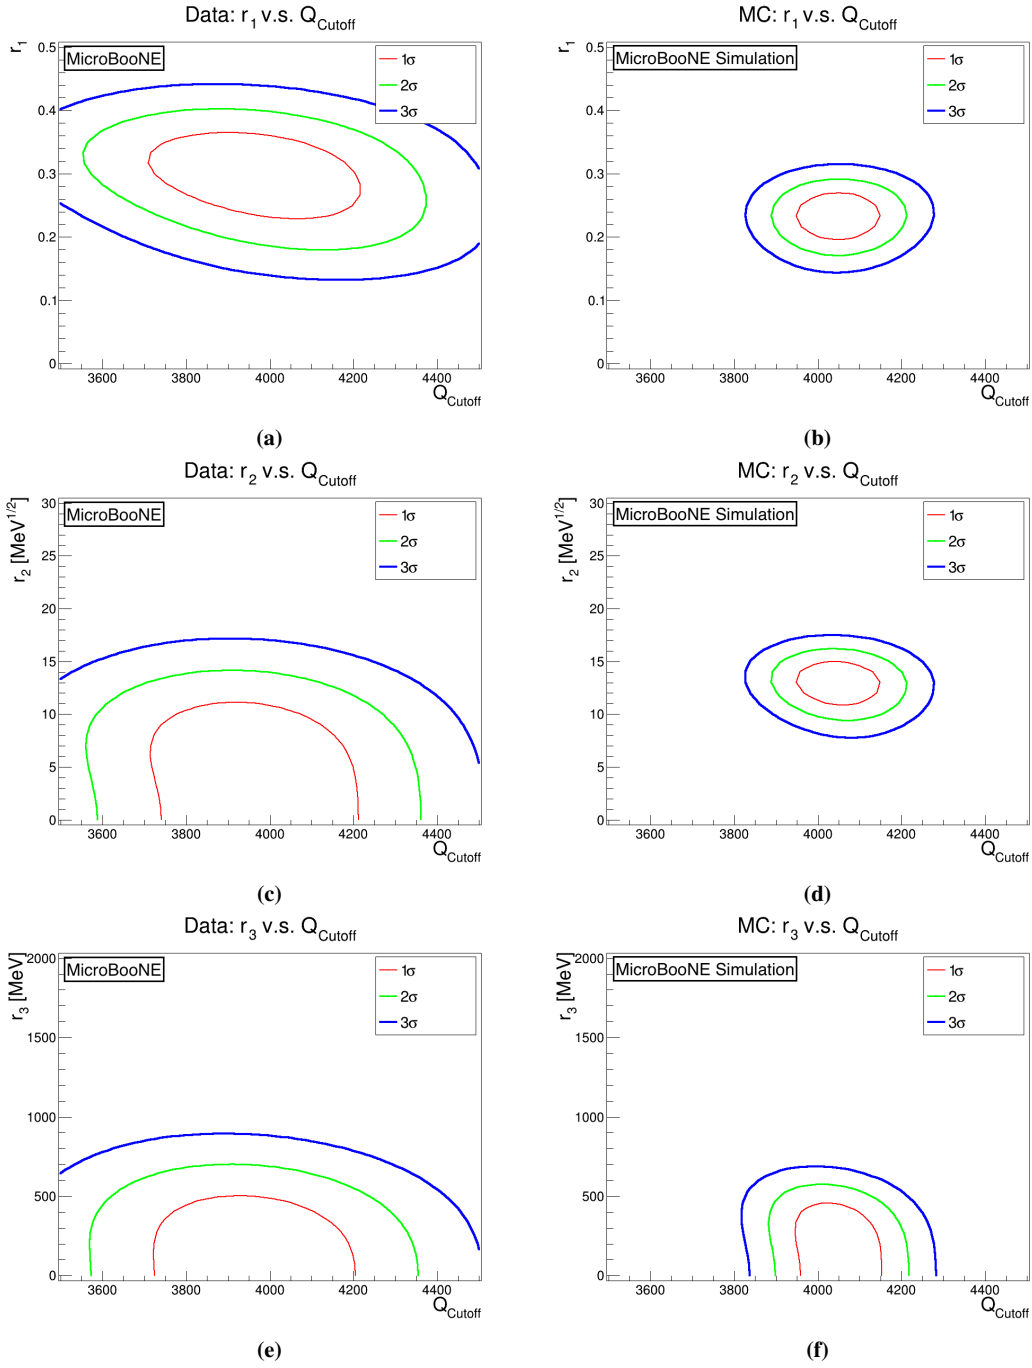


Figure 14. Two-dimensional confidence regions for each resolution parameter in eq. (5.1) vs. Q_{Cutoff} . 1σ , 2σ , and 3σ regions are shown by red, green, and blue curves, respectively. The confidence regions for (a), (c), and (e) come from the fit to data while those in (b), (d), and (f) come from the fit to simulation (MC). The units of each parameter in the plots are as follows: Q_{Cutoff} [Q counts], r_1 [dimensionless], r_2 [Q counts] $^{1/2}$, r_3 [Q counts].




Figure 15. χ^2 from (5.3) as a function of Q_{cutoff} , for data (a) and MC simulation (b). This sample corresponds to $\approx 5.3 \times 10^{19}$ POT. The 1σ allowed regions from Wilks' theorem are shown in shaded regions below each curve.

6 Combined validation of reconstructed shower energy

Both the data/simulation agreement of the shower reconstruction and the absolute scale of the ν_e simulation-derived $Q_{\text{sh-to-MeV}}$ conversion are validated using our two samples by utilizing the π^0 invariant mass of ≈ 135 MeV and the Michel electron spectrum cut-off at ≈ 52.8 MeV. As described in sections 4.2 and 5.2, we have obtained comparison points separately for data and simulation in each sample. These points are shown with statistical uncertainties in figure 16. The $Q_{\text{sh-to-MeV}}$ conversion factor or m values found in section 3 for electrons, leading photons, and sub-leading photons are shown by shaded bands in figure 16. In principle, one expects agreement between the points and the electron and leading photon calibration line. The 1σ ranges in the m value from both the data/simulation Michel cutoff study and the data/simulation π^0 mass study agree well with the best-fit m values from simulated electrons and leading photons. Agreement with the sub-leading photon line is not necessarily expected because of the reconstruction failure cases discussed previously.

Table 5 shows the agreement of data and MC simulation for each point, as well as the agreement of each data and simulation point to the electron best fit value from eq. (3.1). It is seen here that the best fit m values agree between data and simulation for each sample. This validates the use of the same simulation-derived $Q_{\text{sh-to-MeV}}$ conversion value (m_{e^-}) for both data and simulation. While the best fit m values for each sample do not exactly match m_{e^-} within statistical uncertainty, there are factors that may affect this value. These include: detector response modeling, sub-leading photon reconstruction in the π^0 sample, and backgrounds in the Michel e^- sample. Therefore, the 2-6% difference gives an estimate of the scale of the possible data to simulation bias on the shower energy reconstruction. This size effect is acceptable for use in the DL LEE investigation. The $< 6.5\%$ difference between each ($m_{\text{MC}}, m_{\text{Data}}$) and m_{e^-} gives the scale of the detector systematic uncertainty in this reconstruction process.

In addition, as shown in figure 7 and figure 12, the showers found in the two samples are at different energy ranges. The assumption has been made in this analysis that the $Q_{\text{sh-to-MeV}}$ factor




Figure 16. The data and MC simulation points from each the π^0 sample and the Michel e^- sample are compared with the $Q_{\text{sh-to-MeV}}$ electron calibration line used in the DL analysis from eq. (3.1). The $Q_{\text{sh-to-MeV}}$ photon calibration lines (eq. (4.1) and eq. (4.2)) are also included for reference. The shaded regions represent the statistical uncertainty of this given calibration line.

Table 5. Data and MC simulation best fit m values from each sample and comparison to the charge-to-energy conversion factor from eq. (3.1) ($Q_{\text{sh-to-MeV}} = m_{e^-} = 1.26 \pm 0.01 \times 10^{-2}$). Uncertainties in ratios are calculated from the 1σ range of each value. The background adjusted values are used for the π^0 sample.

Sample	m_{MC} [MeV/Q]	m_{Data} [MeV/Q]	$m_{\text{Data}}/m_{\text{MC}}$	m_{MC}/m_{e^-}	m_{Data}/m_{e^-}
π^0	$1.230^{+0.004}_{-0.006} \times 10^{-2}$	$1.236^{+0.005}_{-0.006} \times 10^{-2}$	$1.005^{+0.006}_{-0.006}$	$0.984^{+0.009}_{-0.009}$	$0.979^{+0.008}_{-0.009}$
Michel e^-	$1.31^{+0.03}_{-0.02} \times 10^{-2}$	$1.34^{+0.06}_{-0.06} \times 10^{-2}$	$1.025^{+0.051}_{-0.049}$	$1.038^{+0.022}_{-0.024}$	$1.064^{+0.048}_{-0.048}$

does not change with shower energy. The samples cover the range of values of interest for the DL 1e1p analysis. We conclude that the $Q_{\text{sh-to-MeV}}$ value given in eq. (3.1) is valid for EM showers in both data and simulation at the energy ranges and precision relevant for the MicroBooNE LEE search.

7 Conclusions

This article has reported the updated method of electromagnetic shower reconstruction for the MicroBooNE DL-based LEE analysis. Two samples that allow us to validate our shower reconstruction have been presented: photons produced by the decay π^0 s, and Michel electrons produced when a

ν_μ -sourced stopped muon decays. The reconstruction and selection for each sample was described. The samples show good data/simulation agreement. The shower energy calculation uses a MC ν_e simulation-derived Q_{sh} -to-MeV conversion factor. The absolute scale of the conversion factor, as well as its application to EM showers in both data and simulations, is validated using the π^0 invariant mass of ≈ 135 MeV and the Michel electron cut-off at ≈ 53 MeV. Excellent data/simulation agreement is seen in this study. The Q_{sh} -to-MeV conversion value used in the DL-based analysis is shown to be consistent with both of these physical quantities. The results we present here form the foundation for the MicroBooNE LEE DL-based analysis of $1e1p$ events that will be released in the future.

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.

References

- [1] MINIBooNE collaboration, *Updated MiniBooNE neutrino oscillation results with increased data and new background studies*, *Phys. Rev. D* **103** (2021) 052002 [[arXiv:2006.16883](#)].
- [2] MicroBooNE collaboration, *Search for an anomalous excess of charged-current quasi-elastic ν_e interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction*, [arXiv:2110.14080](#).
- [3] J.S. Moon, *Using Deep Learning Techniques to Search for the MiniBooNE Low Energy Excess in MicroBooNE with $> 3\sigma$ Sensitivity*, Ph.D. thesis, MIT, 2020. [arXiv:2010.14505](#). 10.2172/1767032.
- [4] MicroBooNE collaboration, *Vertex-finding and reconstruction of contained two-track neutrino events in the MicroBooNE detector*, *2021 JINST* **16** P02017 [[arXiv:2002.09375](#)].
- [5] MicroBooNE collaboration, *Design and Construction of the MicroBooNE Detector*, *2017 JINST* **12** P02017 [[arXiv:1612.05824](#)].
- [6] MicroBooNE collaboration, *Measurement of space charge effects in the MicroBooNE LArTPC using cosmic muons*, *2020 JINST* **15** P12037 [[arXiv:2008.09765](#)].
- [7] MicroBooNE collaboration, *Cosmic Ray Background Rejection with Wire-Cell LArTPC Event Reconstruction in the MicroBooNE Detector*, *Phys. Rev. Applied* **15** (2021) 064071 [[arXiv:2101.05076](#)].
- [8] MicroBooNE collaboration, *Neutrino event selection in the MicroBooNE liquid argon time projection chamber using Wire-Cell 3D imaging, clustering, and charge-light matching*, *2021 JINST* **16** P06043 [[arXiv:2011.01375](#)].

- [9] MicroBooNE collaboration, *Noise Characterization and Filtering in the MicroBooNE Liquid Argon TPC*, 2017 *JINST* **12** P08003 [[arXiv:1705.07341](#)].
- [10] MicroBooNE collaboration, *Ionization electron signal processing in single phase LArTPCs. Part I. Algorithm Description and quantitative evaluation with MicroBooNE simulation*, 2018 *JINST* **13** P07006 [[arXiv:1802.08709](#)].
- [11] MicroBooNE collaboration, *Ionization electron signal processing in single phase LArTPCs. Part II. Data/simulation comparison and performance in MicroBooNE*, 2018 *JINST* **13** P07007 [[arXiv:1804.02583](#)].
- [12] MicroBooNE collaboration, *Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE*, *Phys. Rev. D* **103** (2021) 052012 [[arXiv:2012.08513](#)].
- [13] MicroBooNE collaboration, *Reconstruction and Measurement of $O(100)$ MeV Energy Electromagnetic Activity from $\pi^0 \rightarrow \gamma\gamma$ Decays in the MicroBooNE LArTPC*, 2020 *JINST* **15** P02007 [[arXiv:1910.02166](#)].
- [14] MicroBooNE collaboration, *Calibration of the charge and energy loss per unit length of the MicroBooNE liquid argon time projection chamber using muons and protons*, 2020 *JINST* **15** P03022 [[arXiv:1907.11736](#)].
- [15] X. Ji, W. Gu, X. Qian, H. Wei and C. Zhang, *Combined Neyman–Pearson chi-square: An improved approximation to the Poisson-likelihood chi-square*, *Nucl. Instrum. Meth. A* **961** (2020) 163677 [[arXiv:1903.07185](#)].
- [16] S.S. Wilks, *The large-sample distribution of the likelihood ratio for testing composite hypotheses*, *Ann. Math. Stat.* **9** (1938) 60.
- [17] MicroBooNE collaboration, *Michel Electron Reconstruction Using Cosmic-Ray Data from the MicroBooNE LArTPC*, 2017 *JINST* **12** P09014 [[arXiv:1704.02927](#)].
- [18] C. Bouchiat and L. Michel, *Theory of μ -Meson Decay with the Hypothesis of Nonconservation of Parity*, *Phys. Rev.* **106** (1957) 170.

4.5 $1e1p$ Event Selection

The signal events for this analysis, $1e1p$ events consistent with two-body CCQE scattering, are isolated using a series of cuts on kinematic, topological, and MPID-based variables as well as an ensemble of boosted decision trees (BDTs). The $1e1p$ signal selection represents one of my main projects within the MicroBooNE experiment. This section covers in detail the requirements which define the $1e1p$ sample. The $1\mu1p$ control sample is isolated analogously to the $1e1p$ sample and will not be covered in this thesis, though more information can be found in Ref. [148].

4.5.1 Basic Data Selection Criteria

We first apply a series of basic “precuts” meant to remove poorly reconstructed events and obvious backgrounds. We begin with the collection of “track-shower” vertices selected by the algorithm described in section 4.4.2, which represent candidate $1e1p$ events. We remove any vertices that are within 10 cm of the TPC edges or within a region of unresponsive wires between $700 < z \text{ [cm]} < 740$. This cut defines the fiducial volume. We also require that the minimal distance of either the track or the shower to the edge of the TPC is greater than 15 cm. In the case that both prongs are reconstructed close to an edge, we instead require that the distance of each is greater than 5 cm, which helps retain neutrino interactions that begin near the edge of the fiducial volume. Events in the $1e1p$ are further required to be outside of an inefficient region of the U-plane.

We also employ a cut requiring the reconstructed shower energy of the candidate electron to be consistent between the three planes. Even though we rely on the collection plane for the official shower energy estimation, this cut is helpful in removing events that would not lead to a consistent 3D picture of the shower. An example of such an event is shown in figure 4-9. The cut specifically employed in this analysis is

$$\frac{\sqrt{(E_e^U - E_e^V)^2 + (E_e^U - E_e^Y)^2 + (E_e^V - E_e^Y)^2}}{E_e^Y} < 2, \quad (4.7)$$




Figure 4-9: An example of an event that fails the shower energy consistency cut, because the EM shower passes through an unresponsive region of the collection plane.

where $E_e^{[U,V,Y]}$ denote the reconstructed electron candidate energy on the U , V , and Y plane, respectively. Figure 4-10a shows the distribution of the fractional consistency variable (left-hand side of equation (4.7)) for events above and below a $1e1p$ BDT score cutoff of 0.7, where the BDT here is an older iteration of the final BDT ensemble used in this analysis. As can be seen, signal-like events (with a higher BDT score) tend to have lower fractional consistency values compared to background-like events (with a lower BDT score). The upper bound of 2 in equation (4.7) was chosen to retain $\sim 95\%$ of signal-like events while rejecting $\sim 20\%$ of background-like events, as shown in figure 4-10b.

Next, we make cuts on the reconstructed neutrino energy ($200 < E_\nu$ [MeV] < 1200), electron kinetic energy ($K_e > 35$ MeV), and proton kinetic energy ($K_p > 50$ MeV). We also require a forward-going proton and an opening angle between the electron and proton greater than 0.5 rad. Finally, we require that the event can be boosted to the rest frame of the nucleon using the p_{fermi}^T and p_{fermi}^z variables defined in table 4.1.

4.5.2 Boosted Decision Tree Ensemble

The next stage of the $1e1p$ selection uses an ensemble of BDTs to identify signal-like events in MicroBooNE. This is the most powerful cut used in the $1e1p$ analysis. The ensemble consists of 20 BDTs each trained using the XGBoost gradient-boosting algorithm [199]. Each BDT takes as input a set of 19 kinematic variables, such as the reconstructed energies and angles of the electron and proton, and 4 topological variables related to the observed charge, such as the fraction of SparseSSNet-labeled




Figure 4-10: Figure 4-10a shows the distribution of the fractional shower energy consistency variable for events with an old $1e1p$ BDT score above/below 0.7. Figure 4-10b shows the efficiency with which events above/below the old $1e1p$ BDT cutoff pass the fractional consistency cut as a function of the chosen upper bound.

shower pixels in the event. The full suite of variables used to train the $1e1p$ and $1\mu1p$ BDTs is shown in table 4.2. Each BDT itself consists of a collection of regression trees which each assign a continuous score to an event based on a set of decision rules developed using the input variables; the assigned BDT score of the event is given by the sum over the collection [199]. BDTs are trained via gradient-based optimization of an objective function that includes a regularization term to penalize overfitting.

A different ensemble is generated for each of the three run periods, as the detector response changed slightly between them. Thus, there are 60 total BDTs trained for the $1e1p$ analysis. The BDTs are trained using the MicroBooNE MC samples, including the nominal (ν_μ -dominated) BNB sample as well as dedicated intrinsic ν_e , π^0 , and cosmic samples. The signal for the training sample is defined as true ν_e CCQE interactions with one electron and one proton in the final state. We further restrict to well-reconstructed events for which the vertex is reconstructed within 5 cm of the true vertex and the neutrino energy within 20% of the true value. Only non- ν_e CCQE $1e1p$ events are explicitly classified as background. Each BDT in an ensemble is trained using a randomly-selected half of the MC sample for the given run period. Figure 4-11 shows the F score distribution for each variable in one of the $1e1p$ BDTs from the ensemble for each run period. The F score represents the number of times each variable is used across all trees in the BDT; thus, it can be thought of as a

measure of the importance of each variable.

The $1e1p$ score for an event is evaluated by calculating the average BDT score across the ensemble. This reduces the dependence of the ensemble score on the specific training sample of any single BDT. Ensemble-based methods have been shown to reduce variance on the output likelihood variable compared to a single classifier [200]. When computing the average score for events in the MicroBooNE MC sample, we omit BDTs for which the event appeared in that BDT's training sample. It is extremely unlikely ($p = 2^{-20}$) for an event to appear in the training sample of all BDTs in an ensemble; thus, we can use the BDT ensemble to isolate a $1e1p$ signal sample prediction from the MicroBooNE MC without throwing away any events. The score is normalized to the range $[0,1]$, where higher scores indicate a higher signal likelihood. We show the distribution of the ensemble-averaged $1e1p$ BDT score in figure 4-12, which shows that true ν_e CCQE events tend to peak toward 1 while all other events peak toward 0. Signal events are required to have an average score greater than 0.95, which was chosen to optimize sensitivity to the eLEE model described in section 4.2.

We have performed a series of tests of the robustness of the BDT ensemble method. We first examined the signal selection power of a BDT ensemble trained on simulated events from a given run period when used to infer the signal likelihood of simulated events from a different run period. This is useful for two reasons. For one, it provides a method for testing the ensemble on events that were in neither the training nor the validation sample of the constituent BDTs. Testing on a sample independent of the validation sample is important because the BDT training was halted when the classification error on the validation set didn't improve after 50 training iterations. It also tests the impact of removing training events from the MC prediction. When using the BDT ensemble of a given run period to evaluate the signal likelihood of events from the same run period, one must remove events that appeared in the training sample of each constituent BDT. In this study we use the BDT ensemble trained on run period 2 to evaluate simulated ν_e events from run period 3 and vice versa, so one does not need to remove training events.

Figure 4-13 shows the results of this run period swap study. As one can see, the

differences between the selection using the correct and incorrect run period ensemble are small. The correct run period performs slightly better, as expected due to small differences in the detector status specific to each run period. However, the relatively small change in ν_e selection efficiency suggests that the BDT ensemble is able to perform well on simulated events that do not appear in either the training or validation sample of constituent BDTs. One can interpret the loss in performance shown in figure 4-13 as a bound on the potential impact of applying the BDT ensemble to events that are not in the training or validation samples of the entire ensemble. The drop in performance is $< 10\%$ over most of the energy range—an effect which is smaller than our systematic error before the ν_μ constraint described in chapter 5 ($\sim 15\%$) and much smaller than the statistical error in each reconstructed neutrino energy bin after the final selection ($\sim 50\%$).

Next, we investigated the impact of using only a subset of the BDTs in the ensemble to calculate the average BDT score of an event. This is relevant because we omit BDTs that contained a given MC event in the training sample when calculating the average BDT score for that event. However, this analysis uses all 20 BDTs in the ensemble to calculate the average BDT score of events in the data. Therefore, to verify that comparisons between data and simulation are robust, one needs to ensure that removing BDTs from the ensemble does not significantly bias the average score calculation.

To this end, consider S_n to be the $1e1p$ BDT ensemble average score after removing n BDTs from the ensemble. For events in data, the BDT score is S_0 while for simulated events, the BDT score is S_n for some $n \in \{1, \dots, 20\}$. Figure 4-14 shows the fractional difference $(S_n - S_0)/S_0$ as a function of the number of omitted BDTs n over signal-like (BDT score > 0.95) events in the simulation from run period 2 and run period 3, respectively. As in the previous study, in order to avoid bias from BDT training we use the BDT ensemble trained on run period 2 to evaluate the signal likelihood of simulation events from run period 3 and vice versa. The data points and error bars in figure 4-14 indicate the average and standard deviation of the fractional difference over the simulation sample, respectively. The red histograms show the actual distribution

of omitted BDTs over the simulation from each run period. One can see that the average BDT score does not exhibit significant bias upon removing BDTs from the calculation. Also, the standard deviation of the fractional difference only becomes larger than 1% when removing $\gtrsim 17$ BDTs from the calculation. This only happens for $\ll 1\%$ of simulated events, as expected.

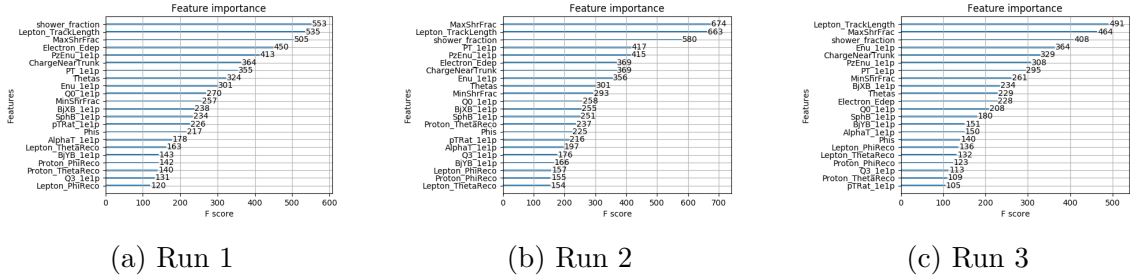


Figure 4-11: The F score of each variable for one of the $1e1p$ BDTs in the ensemble from run 1, run2, and run 3.

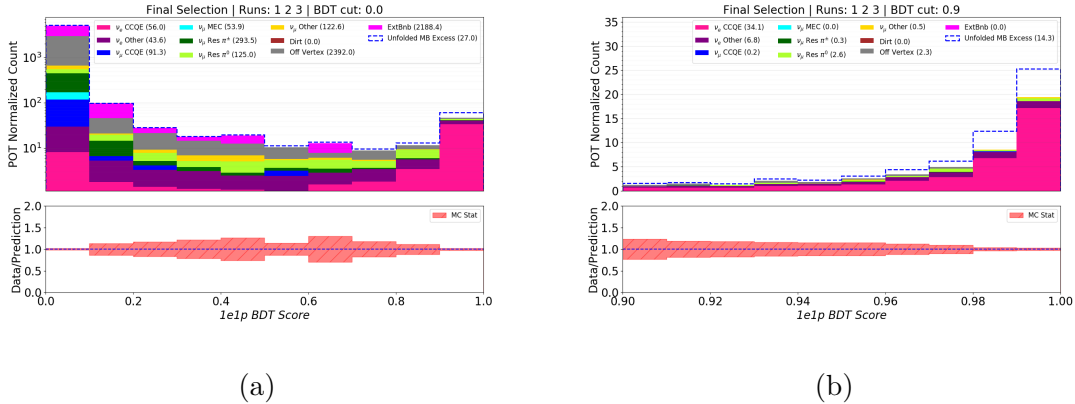




Figure 4-12: Figure 4-12a and figure 4-12b show the MC distribution of the $1e1p$ BDT ensemble average score over all three run periods for the full $[0,1]$ range and zoomed in to the $[0.95,1]$ range, respectively.

4.5.3 Particle Identification Cuts

Finally, we employ a series of particle-identification requirements to clean up the remaining backgrounds that survive the $1e1p$ BDT ensemble cut. The first of these is a cut on the invariant π^0 mass. If the shower reconstruction algorithm described in section 4.4.3 is able to identify a second EM shower, the reconstructed M_{π^0} must




(a)




(b)

Figure 4-13: Figure 4-13a and figure 4-13b show the predicted ν_e event rate in run period 2 and run period 3, respectively, using both the run period 2 ensemble and the run period 3 ensemble.



(a)



(b)

Figure 4-14: Figure 4-14a and figure 4-14b show the fractional difference in average BDT score $(S_n - S_0)/S_0$ as a function of the number of omitted BDTs n over the simulation from run period 2 and run period 3, respectively. The red histogram shows the actual distribution of the number of omitted BDTs over the run period 2 and run period 3 simulation samples, respectively. Scores are calculated using the run period 3 and run period 2 BDT ensemble, respectively.

be less than 50 MeV. Another cut requires the ratio of the MPID γ and e^- image scores to be less than 2. This helps remove remaining π^0 events that survive the M_{π^0} cut. We also require the MPID muon interaction score to be less than 0.2, which helps remove ν_μ $CC\pi^0$ events where the muon gets mistaken for a proton. This last cut is only applied for events in which $E_{e^-} > 100$ MeV, as MPID tends to assign low electron scores and high muon scores to true electrons below 100 MeV, as shown in figure 4-15. This is because electrons stop radiating below this energy, becoming more track-like and thus difficult to distinguish from muons. The specific values used for

the cuts described here were chosen by optimizing the signal-to-background ratio in the $1e1p$ sample. Finally, if an event has more than one reconstructed neutrino vertex candidate passing all selection criteria, we keep only the vertex with the highest $1e1p$ ensemble BDT score.




Figure 4-15: Figure 4-15a and figure 4-15b show the MPID electron and muon score, respectively, as a function of the reconstructed electron energy in intrinsic ν_e MC events.

4.5.4 The Final $1e1p$ Sample

The requirements described throughout this section define the $1e1p$ sample; they are summarized and compared with those in the $1\mu1p$ sample in table 4.3. This results in a highly-pure $1e1p$ signal sample—75% of events are true ν_e CCQE interactions. Figure 4-16 shows the predicted E_ν^{range} distribution for this sample, including the prediction from the eLEE model. This figure also shows the systematic uncertainty and non- ν_e background prediction, which will be discussed further in chapter 5. The neutrino energy distribution in figure 4-16 forms the basis of the statistical results discussed in chapter 5.

The neutrino energy resolution is 16.5% for selected $1e1p$ events, as shown in figure 4-17. This figure also indicates a slightly higher rate of events with under-predicted neutrino energy, owing mainly to showers and tracks that pass through unresponsive regions of the detector. Figure 4-18 reports the signal efficiency and total number of ν_e CCQE events in this sample after applying each of the three sets

of cuts. The efficiency peaks at lower neutrino energies ~ 300 MeV, which is optimal for testing the eLEE model. Note that these efficiencies in figure 4-18a are reported with respect to all identified vertices. If we include the effect of the vertex algorithm, the average efficiency for retaining signal events in the $1e1p$ sample is 6.6% [148]. This is by design—in choosing a very exclusive signal definition, we optimize for a high-purity sample at the price of lower efficiency.




Figure 4-16: The E_ν^{range} distribution for the $1e1p$ signal sample, showing only the predicted event rate from the MC. The prediction from the eLEE model is shown in the dashed blue line.






Figure 4-17: Figure 4-17a shows the distribution of the fractional error on the neutrino energy for MC events in the $1e1p$ signal sample, restricted to $200 < E_\nu^{\text{Range}} [\text{MeV}] < 1200$. Figure 4-17b shows the 2D distribution of fractional error as a function of the true neutrino energy.



(a)



(b)

Figure 4-18: Figure 4-18a shows the post-vertex-identification efficiency of true ν_e CCQE selection for subsequent stages of the $1e1p$ cuts. Figure 4-18b shows the true ν_e CCQE event rates over the full run 1-3 dataset after subsequent stages of the $1e1p$ cuts.

Variable Name	Definition
Base Variables	
K_p K_ℓ M_ℓ, M_n, M_p $\cos\theta_p, \cos\theta_\ell$ ϕ_p, ϕ_ℓ $P_p = (E_p, p_p), P_\ell = (E_\ell, p_\ell)$ B	Kinetic energy of proton determined from range [24] Kinetic energy determined from range for muons [24] and from calorimetry for electrons [172] Masses of the lepton, neutron and proton $p_p^z/p_p, p_\ell^z/p_\ell$ $atan2(p_p^y, p_p^x), atan2(p_\ell^y, p_\ell^x)$ Reconstructed 4-vector of the proton, lepton Binding Energy for argon; the analysis assumes $B = 40$ MeV
Definitions Related to Neutrino Energy	
E_ν^{range} (Default value of E_ν) E_ν^{QE-p} $E_\nu^{QE-\ell}$	$K_p + K_\ell + M_\ell + M_p - (M_n - B)$ $\left(\frac{1}{2}\right) \frac{2 \cdot (M_n - B) \cdot E_p - ((M_n - B)^2 + M_p^2 - M_\ell^2)}{(M_n - B) - E_p + \sqrt{(E_p^2 - M_p^2)} \cdot \cos\theta_p}$ $\left(\frac{1}{2}\right) \frac{2 \cdot (M_n - B) \cdot E_\ell - ((M_n - B)^2 + M_\ell^2 - M_p^2)}{(M_n - B) - E_\ell + \sqrt{(E_\ell^2 - M_\ell^2)} \cdot \cos\theta_\ell}$
Δ^{QE} (2-Body Consistency)	$\sqrt{(E_\nu^{range} - E_\nu^{QE-p})^2 + (E_\nu^{range} - E_\nu^{QE-\ell})^2 + (E_\nu^{QE-\ell} - E_\nu^{QE-p})^2}$
Event Kinematics	
Q^2 Hadronic Energy (E_{had}) Björken's Scaling x (x_{Bj}) Björken's Scaling y (y_{Bj}) Opening angle	$2E_\nu^{range}(E_\ell - P_\ell^z) - M_\ell^2$ $E_\nu^{range} - E_\ell$ $Q^2/2M_n E_{had}$ E_{had}/E_ν^{range} $\cos^{-1}(\hat{p}_\ell \cdot \hat{p}_p)$
p_T p_L α_T ϕ_T	$\sqrt{(p_\ell^x + p_p^x)^2 + (p_\ell^y + p_p^y)^2}$ $p_p^z + p_\ell^z$ $\cos^{-1}\left(-\frac{\vec{P}_T^l \cdot \vec{P}_T}{ \vec{P}_T^l \vec{P}_T }\right)$ $\cos^{-1}\left(-\frac{\vec{P}_T^l \cdot \vec{P}_T^p}{ \vec{P}_T^l \vec{P}_T^p }\right)$
Boosting Parameters	
p_{Fermi}^T p_{Fermi}^z	$p_p^T + p_\ell^T$ $p_p^z + p_\ell^z - E_\nu^{range}$

Table 4.1: The definition of kinematic variables used throughout the two-body CCQE analysis.

Variable	Used in $1\mu 1p$ BDT	Used in $1e1p$ BDT
Variables Used in BDTs, Based on Ionization		
Charge within 5 cm of vertex	Yes	Yes
Shower charge in event image / shower charge clustered as electron	No	Yes
Proton shower fraction	No	Yes
Electron shower fraction	No	Yes
Variables Used in BDTs, Related to Energy Measurements		
Neutrino Energy	Yes	Yes
Energy of electromagnetic shower	No	Yes
Lepton length	Yes	Yes
Proton length	No	Yes
$p_z - E_\nu$	No	Yes
Variables Used in BDTs, Related to 2-Body Scattering Consistency		
Bjorken's x	Yes *	Yes *
Bjorken's y	Yes *	Yes *
QE Consistency	Yes *	Yes *
Q_0	Yes	Yes
Q_3	Yes	Yes
Variables Used in BDTs, Related to Transverse Momentum		
α_T	Yes	Yes
Event p_T	Yes	Yes
Event p_T/p ("PTrat")	Yes	Yes
ϕ_T	Yes	No
Variables Used in BDTs, Related to Angles		
Proton ϕ	Yes	Yes
Proton θ	Yes	Yes
Lepton ϕ	Yes	Yes
Lepton θ	Yes	Yes
$\phi_p - \phi_\ell$	Yes	Yes
$\theta_p + \theta_e$	No	Yes
Variables Useful for Comparison, Not Used in Either BDT		
η (Normalized average ionization difference)	No	No
Opening Angle	No	No
x Vertex	No	No
y Vertex	No	No
z Vertex	No	No

Table 4.2: The suite of variables used to isolate and analyze the $1e1p$ and $1\mu 1p$ samples. Variables used in the BDT ensemble for each sample are specified. The "*" character indicates that the variable is calculated in the rest frame of the struck nucleon. The mathematical definitions of many of these variables appear in Table 4.1.

Name/Description of Code	Requirement for $1\mu 1p$	Requirement for $1e1p$
Cuts/Selections Applied in External Code		
Common Optical Filter	> 20 PMT hits w/i 6 ticks in beam window $\& \leq 20$ hits w/i 6 ticks in $2 \mu s$ prior to spill	Same
WC tagger	Mask cosmic-associated charge	Same
Preselection Cuts for Event Selection		
Run Quality Requirements: Good Run Flag	True	Same
Vertex Requirements: $1\ell 1p$ consistent Edge fiducial No dead region	Two prongs of at least 5cm length, no additional prongs > 10 cm from active volume edge outside $z = 700$ to 740 cm	Same Same Same
Containment Requirements: Edge containment Efficient region	closest approach > 15 cm from edge N/A	Same outside $y = (1/\sqrt{3}) * z - 117$ to $y = 1/\sqrt{3} * z - 80$ cm.
Particle Energy Requirements: Neutrino energy (range-based) Lepton kinetic energy Proton kinetic energy	$E_\nu < 1200$ MeV $K_\ell > 35$ MeV $K_p > 50$ MeV	Same Same Same
Analysis Orthogonality Cut	Max Shower Frac < 0.2	Max Shower Frac > 0.2
Other Basic Quality Requirements: Opening Angle Proton θ Boostable Shower 3-view energy consistency Vertex with BDT score that is	> 0.5 radians $\cos(\theta_p) > 0$ $\gamma > 1, 0 < \beta < 1$ N/A lowest (<i>i.e.</i> best)	Same Same Same Inconsistency $< 200\%$ highest (<i>i.e.</i> best)
BDT Requirements for Event Selection		
BDT requirement	ν_μ BDT < 0.4	ν_e BDT > 0.95
PID Requirements for Event Selection		
π^0 mass muon MPID interaction score γ/e image score ratio proton MPID interaction (max of 3 planes)	N/A N/A N/A > 0.9 if $E_\nu < 400$ MeV	π^0 mass < 50 MeV < 0.2 if $E_e > 100$ MeV < 2 N/A

Table 4.3: The specific used to define the $1\mu 1p$ and $1e1p$ samples. For definitions of kinematic variables, see Table 4.1.

Chapter 5

The MicroBooNE Electron Neutrino Analysis: Results and Discussion

This chapter presents the first results from the two-body CCQE analysis described in chapter 4. They come from data taken during the first three MicroBooNE run periods, corresponding to 6.67×10^{20} POT after data quality cuts. The signal data were examined for the first time in Summer 2021 and presented to the community in October 2021 [148]. The results from the two-body CCQE analysis were accompanied by results from the inclusive [170] and MiniBooNE-like [171] ν_e analyses. All three MicroBooNE ν_e analyses are summarized in Ref. [116].

This chapter focuses on the statistical results from the ν_e two-body CCQE analysis. We cover the background prediction, systematic error evaluation, ν_μ $1\mu 1p$ constraint procedure, and blinded analysis approach in detail. We then discuss the series of statistical tests performed using our $1e1p$ signal sample. We close with a discussion of the MicroBooNE results, including a presentation of two non-MicroBooNE publications examining the implications of the MicroBooNE data under two different explanations of the MiniBooNE LEE: $3 + 1$ oscillations and excess $\bar{\nu}_e$ interactions.

Publications covered in this chapter for which I either held a leading role or made major contributions: [102, 116, 148, 201]

5.1 First Results from the Two-Body CCQE Analysis

The two-body CCQE analysis observed 25 events across data from the first three run periods. This is in good agreement with the total prediction in the 200 – 1200 MeV range, which is $27.4 \pm 3.8_{\text{sys}} \pm 5.2_{\text{stat}}$ before the $1\mu 1p$ constraint discussed in section 5.1.3 [148]. One of these events is shown in figure 5-1, including both pixel intensity images from the TPC readout and SparseSSNet pixel-labeled images on all three planes. Event displays for the remaining 24 selected events are included in Appendix B. The E_ν^{range} distribution of these events is shown in figure 5-2, where the prediction is given both in terms of the underlying interaction type and the final state event topology. The $1e1p$ BDT ensemble average score distribution is shown in figure 5-3. Figure 5-4 shows the distributions of four more interesting variables: E_e , E_p , θ_e , and $E_\nu^{QE-\ell}$. A full suite of data-prediction comparisons in 36 variables is provided in the Supplemental Material for Ref. [148], including those described in table 4.1 and those used to train the BDT ensembles.

One can see clearly in figure 5-2 that no excess of ν_e $1e1p$ candidate events is observed at low energy in the two-body CCQE analysis. Instead, we observe a slight deficit compared to even the nominal prediction at low energies, in conflict with the eLEE model represented by the blue dashed line. This is a trend that is observed across all three ν_e analyses in MicroBooNE [116]. The extent to which we agree or disagree with the eLEE model will be discussed further in section 5.2.

One might also notice an apparent excess of data events in the 800-900 MeV bin in figure 5-2. This appears to be a statistical fluctuation; other kinematic variables do not exhibit such an excess of events, including the closely related $E_\nu^{QE-\ell}$ distribution shown in figure 5-4d. If this were a systematic effect rather than a fluctuation, one would expect similar behavior in the distributions of other variables. While these statements will be quantified further in section 5.2, it is useful to include such a qualitative discussion at this point.

The following subsections discuss the details behind our prediction in the E_ν

distribution. It is important to note that while we have already shown the observed data in our signal channel, the procedures outlined in the following subsections were frozen before examining this dataset. Specifically, we followed the blinded analysis approach outlined in section 5.1.4 to avoid biasing our selection.

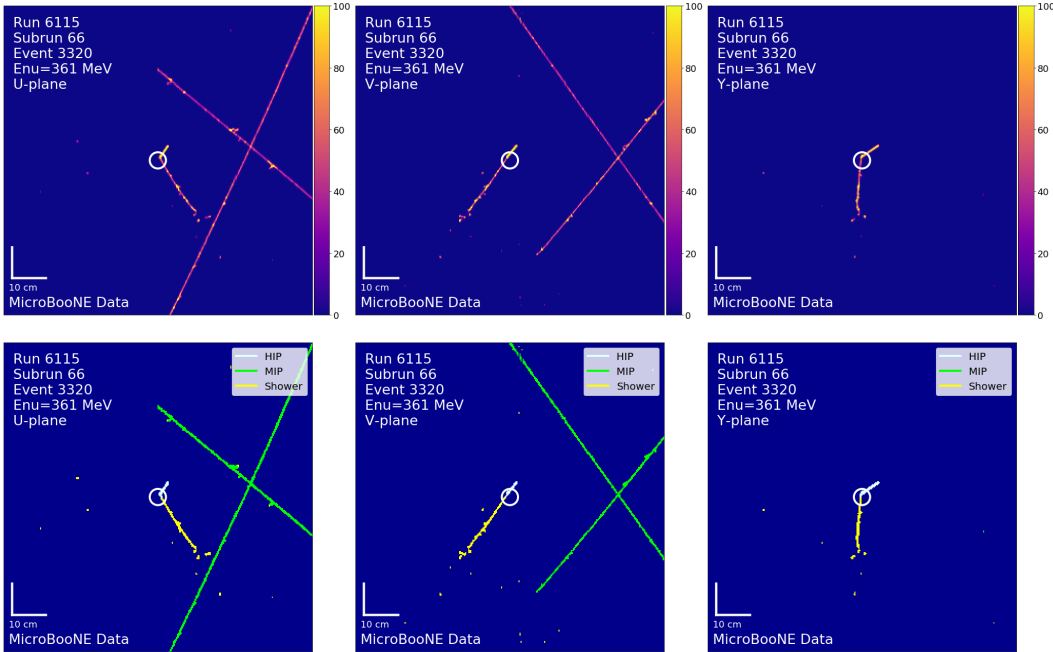


Figure 5-1: Top: pixel intensity (color scale is in PIU as defined in section 4.4); Bottom: SparseSSNet labels; Left to Right: U, V, Y, planes. The white circle indicates the reconstructed vertex. The horizontal axis corresponds to the wire plane direction and the vertical axis corresponds to the electron drift direction, which is measured using the arrival time of charge on the wires.

5.1.1 Background Estimation

In table 5.1 we give a breakdown of the MC events in the “background” category of figures 5-3 and 5-4 by their interaction channel and event topology. The “off-vertex” label in this table refers to events for which the reconstructed neutrino vertex is further than 5 cm away from the true neutrino vertex. These events can happen when a cosmic muon track happens to cross a photon shower from a π^0 decay, or a muon decays to a Michel electron; in both cases, the muon is mistaken for a proton. Table 5.1 indicates that ν_μ resonant π^0 events are the most common background

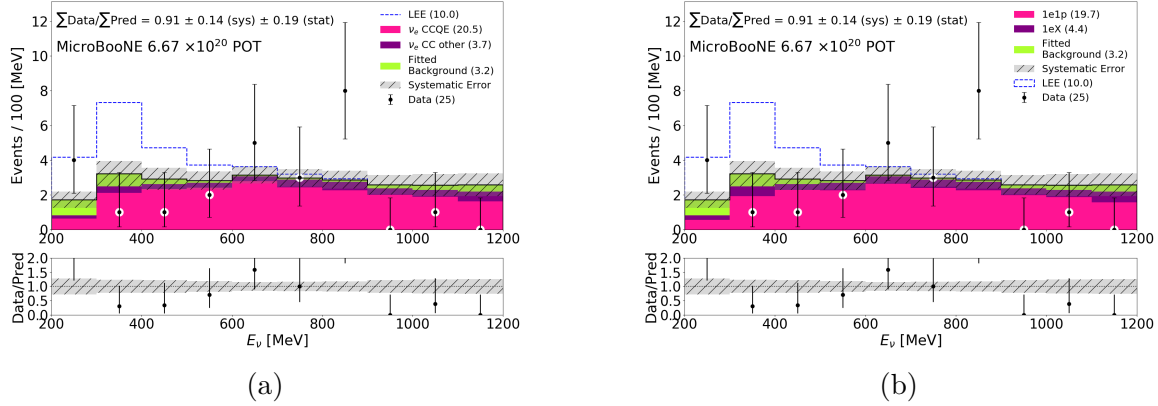


Figure 5-2: The $1e1p$ sample E_ν distribution, comparing data (black points) to the unconstrained prediction (stacked histogram) in the $200 < E_\nu < 1200$ MeV region. The eLEE model prediction is represented by the dashed blue line. The prediction is presented in terms of both interaction type (figure 5-2a) and final state topology (figure 5-2b).




Figure 5-3: Average $1e1p$ BDT ensemble score distribution comparing data to the unconstrained prediction.

interaction channel, followed by off-vertex events. One can also see that events with a π^0 in the final state are the dominant background topology.

There are not many MC events to inform the numbers reported in table 5.1. The tools developed for this analysis produce a highly pure ν_e CCQE sample, resulting in a low-statistics simulation sample for assessing the non- ν_e background to the $1e1p$ signal. This is apparent in the fluctuations of the background prediction in figures 5-3 and 5-4. In order to obtain a more robust prediction of this background, we have elected to leverage information on the energy distribution of non- ν_e background events

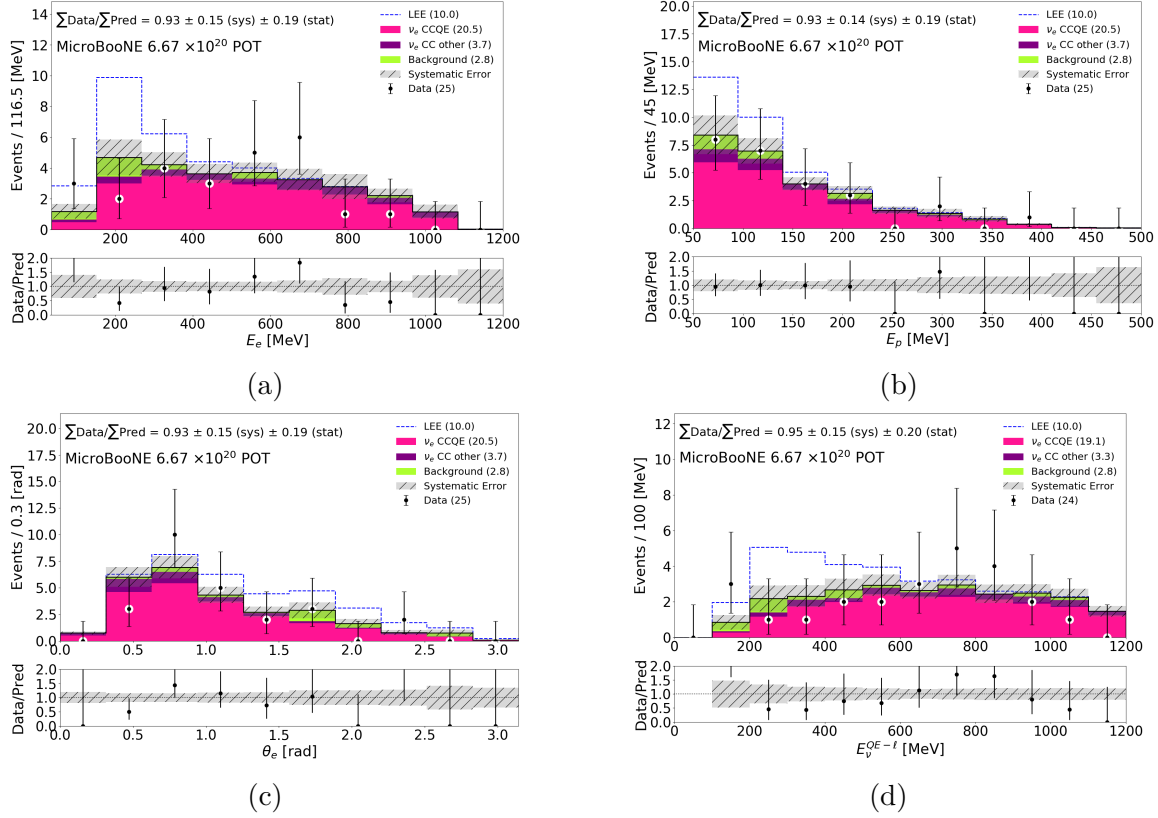


Figure 5-4: Comparison between data and unconstrained prediction in the E_e (figure 5-4a), E_p (figure 5-4b), θ_e (figure 5-4c), and $E_\nu^{QE-\ell}$ (figure 5-4d) distributions of the $1e1p$ sample.

at a loose BDT score cutoff of 0.7 and extrapolate this to our signal cutoff of 0.95. This is accomplished by fitting a parameterized Landau+linear probability density function (PDF) to the background energy distribution at the loose cutoff and scaling this prediction to the signal cutoff. As most of these backgrounds come from ν_μ interactions, this is referred to as the ν_μ background fit.

The Landau+linear shape is motivated empirically by the observation of a rise in the background rate toward the lowest energies ($\lesssim 500$ MeV) and a smaller rise toward the highest energies ($\gtrsim 800$ MeV), both for the loose BDT score cutoff and the signal cutoff, as shown by the blue points in figure 5-5. The Landau portion of the fit is additionally motivated by the observation that a majority of ν_μ background events contain a π^0 in the final state, for which one of the decay photons is misinterpreted as an electron shower. The reconstructed neutrino energy of these events is governed

predominately by the energy of this photon, which has a tail extending to higher energies caused by pions with high momentum in the lab frame. This tail is a characteristic feature of the Landau function. The output of the fit and resulting error are used in place of the raw prediction and statistical error for the ν_μ backgrounds in this analysis. We note here that the simulated ν_μ events with a π^0 in the final state are weighted according to the observation in our dedicated π^0 sample. The method for calculating these weights is described in Ref. [148].

The predicted background spectrum in each reconstructed neutrino energy bin $f(E_i)$ is generated by integrating the Landau+linear PDF $p(E)$ within that bin. Specifically,

$$\begin{aligned}
 p(E) &= \exp\left[\frac{-(E' + e^{-E'})}{2}\right] + aE, \\
 E' &= (E - \mu)/\sigma, \\
 f(E_i) &= \int_{E_i - \delta E_i}^{E_i + \delta E_i} p(E) dE,
 \end{aligned}
 \tag{5.1}$$

where μ and σ are the center and width of the Moyal approximation of the Landau function [202], a is the linear slope parameter, E_i is the center of the i 'th energy bin, and δE_i is half of the bin width. The Landau+linear fit is carried out using only shape information at a loose BDT score cutoff of 0.7. In order to get the overall normalization, we fit the BDT score distribution of the backgrounds to a linear PDF $\tilde{p}(x)$, which we can integrate to get the total expected number of background events $\tilde{f}(x)$ for a given BDT score cutoff x ,

$$\tilde{p}(x) = mx + b \quad \tilde{f}(x) = \int_x^1 \tilde{p}(x) dx,
 \tag{5.2}$$

where m and b are the slope and bias parameters for the BDT score distribution, respectively. The resulting shape+normalization fit for the background distribution at the loose BDT cutoff of 0.7 and signal cutoff of 0.95 are shown in Fig. 5-5. One can see that the fit agrees with the raw MC prediction within statistical error in both cases and that the error on the fit is generally reduced compared to the simulation

statistical error.

The errors on the fits in Fig. 5-5 are obtained by simulating pseudo-experiments according to the covariance matrix of the fit parameters. This is accomplished efficiently using Cholesky decomposition [203]. The overall normalization error on the background rate at a given BDT score cutoff is calculated using the error on the linear fit to the BDT score distribution of background events. For the Landau+linear fit, we generate a shape-only covariance matrix for the ν_μ backgrounds in the nominal reconstructed neutrino energy bins. The overall normalization error has been added as a fully-correlated contribution to each bin of the shape-only E_ν covariance matrix. The uncorrelated uncertainty from the fit, given by the square root of the diagonal entries of the full (shape \oplus normalization) covariance matrix, is indicated by the orange band in Fig. 5-5. This covariance matrix replaces the nominal statistical uncertainty on the simulated background, given by the quadrature sum of weights in each E_ν bin, which is indicated by the shaded blue band in Fig. 5-5.

The performance of the fit is evaluated on data by examining events slightly below the signal region, with a BDT score in the range [0.7, 0.95]. These events are shown in Fig. 5-6. The raw MC prediction is given by the stacked histogram, while the prediction incorporating the ν_μ background fit on top of the simulated ν_e prediction is given by the red line. The data agree well with both predictions, indicating the consistency of the fit method presented here with both the observed and MC-predicted background rate.

5.1.2 Evaluation of Systematic Uncertainties

There are five main sources of systematic uncertainty in this analysis, related to the beam flux, neutrino-nucleus cross section modeling, hadron re-interaction modeling, detector simulation, and finite statistics in the MC-based prediction. These are handled within a covariance matrix formalism. For each binned distribution, the diagonal entries of the covariance matrix contain the variance in each bin, while the off-diagonal entries contain the covariance between different bins. Thus, the covariance matrix is also symmetric. The use of a covariance matrix is essential for encoding correlations

Interaction Channel	Predicted Rate
ν_μ resonant π^0	1.26
ν_μ resonant π^\pm	0.21
ν_μ CCQE	0.14
ν_μ other	0.19
Off-vertex	0.93
Event Topology	Predicted Rate
$1\mu N\pi^0$	0.57
$0\mu N\pi^0$	1.09
$1\mu 1p$	0.14
Off-vertex	0.93

Table 5.1: Breakdown of MC events in the “background” category of figures 5-3 and 5-4 over the range $200 < E_\nu < 1200$ MeV. The events are partitioned both by the interaction channel and the event topology.

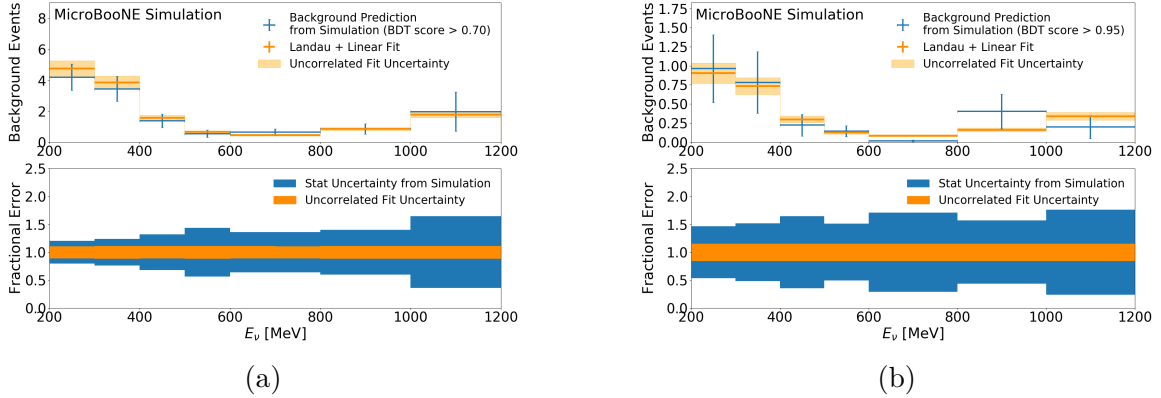


Figure 5-5: The fit to the ν_μ background distribution to the $1e1p$ analysis. The shape fit is performed at a loose BDT score cutoff of 0.7 (figure 5-5a) and scaled to the signal cutoff of 0.95 (figure 5-5b). Blue points represent the prediction from the simulation, with error bars representing the Gaussian approximation of the statistical error (quadrature sum of event weights). The orange line and corresponding shaded region represent prediction and uncertainty, respectively, coming from the Landau+linear fit.

between binned distributions of selected events in the $1e1p$ and $1\mu 1p$ samples, which is relevant for the constraint procedure described in section 5.1.3. The total covariance matrix for a given distribution is simply the element-wise sum of the covariance matrices for each source of systematic uncertainty. A few definitions will be helpful for our discussion here; given a covariance matrix M_{ij} and binned prediction N_i , the fractional covariance matrix is given by $F_{ij} \equiv M_{ij}/(N_i N_j)$ and the correlation matrix




Figure 5-6: The data and MC prediction for events with a $1e1p$ BDT score inside $[0.7, 0.95]$. Good agreement is observed between data and prediction. The prediction incorporating the Landau+linear background fit is shown by the red line.

is given by $\rho_{ij} \equiv M_{ij} / \sqrt{M_{ii}M_{jj}}$. The correlation matrix is meant to capture the degree of correlation between two bins normalized within the range $[-1, 1]$, where diagonal entries are 1 by construction.

The flux, cross section, and re-interaction uncertainties are evaluated using a reweighting method, in which event weights are modified according to tunable parameters. Variations in the distributions of reconstructed variables are used to determine the associated covariance matrix. The detector systematic covariance matrix is evaluated using a dedicated set of simulation samples that each alter the detector response in a different way. These alterations to the detector response are derived using cosmic muon data; the full procedure is described in detail in Ref. [204].

Flux uncertainties are related to hadron production in the target, secondary hadron interactions, and focusing efficiency of the BNB magnetic horn. Neutrino events are reweighted according to the species and momentum of the parent meson [14, 205]. The largest flux uncertainty comes from π^+ production in the target. The flux uncertainties are also highly correlated between the $1e1p$ and $1\mu1p$ samples, as they both come predominately from the same meson decay chain, $\pi^+ \rightarrow \nu_\mu(\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu)$.

The cross section covariance matrix is calculated using both interaction-specific

reweightable parameters and FSI-related reweightable parameters in `GENIE v3.00.06`. Thus, each MC event is reweighted according to its interaction channel and any final state particles which underwent FSI. The full suite of reweightable parameters is described in Ref. [180]. We also evaluate uncertainties related to the difference between the ν_e CCQE and ν_μ CCQE cross sections [206]. The largest cross section uncertainties come from CCQE-related parameters, which is expected given our signal definition, and FSI-related parameters. FSI uncertainties are especially important in this analysis, as FSI can cause ν_e CCQE events to deviate from two-body scattering kinematics and also cause non- ν_e CCQE events to look more consistent with two-body scattering. Both CCQE-related and FSI-related effects impact the $1e1p$ and $1\mu1p$ samples similarly and thus are at least partially addressed by the constraint procedure in section 5.1.3.

Hadron re-interaction uncertainties describe the scattering of final state protons and π^\pm off of argon nuclei. The covariance matrix is calculated by reweighting events according to `Geant4` truth-level information about the hadron trajectory [207]. These uncertainties are a small effect in this analysis.

The detector systematic covariance matrix is calculated using a set of dedicated simulation samples that vary certain aspects of the detector response. These include the amplitude and width of TPC signals as a function of the x and (y, z) positions as well the directions θ_{XZ} and θ_{YZ} of the charged particle; the electron-ion recombination rate; the electric field map in the TPC; and the light yield, attenuation, and Rayleigh scattering length. As these samples are computationally expensive to generate, we are not left with many MC events with which to construct a covariance matrix. Thus, we rely on kernel density estimation [208, 209] to smooth the predictions in each detector variation sample. This procedure does not work for the ν_μ backgrounds to the $1e1p$ analysis since we are especially starved for MC statistics here, as discussed in section 5.1.1. We instead assign a conservative uncorrelated 20% fractional detector systematic uncertainty on these events, based on the extent of differences in the total rate of selected background events in the detector variation samples.

The variance due to finite MC statistics is simply given by the quadrature sum

of event weights in each bin. These uncertainties are uncorrelated and give a diagonal covariance matrix. The exception is finite MC statistics uncertainties on the ν_μ backgrounds in the $1e1p$ analysis, which are calculated using the procedure described in section 5.1.1.

Finally, we note that to evaluate the reweightable systematic uncertainties on the ν_μ backgrounds to the $1e1p$ analysis, we use the selected MC events to calculate a fractional covariance matrix, which is then scaled by the parameterized background prediction derived in section 5.1.1. When performing the constraint in section 5.1.3, we do not consider off-diagonal systematic uncertainties between ν_μ backgrounds in the $1e1p$ sample and bins in the $1\mu1p$ sample, as it is not straightforward to compute this correlation.

The fractional uncertainty in each bin of the E_ν^{range} distribution in the $1e1p$ and $1\mu1p$ samples is shown in figure 5-7. These correspond to the square root of the diagonal entries of the total systematic covariance matrix calculated from all sources described above, divided by the predicted event rate in each bin. In the $1e1p$ sample, flux, cross section, and detector uncertainties dominate at low energies while detector uncertainties dominate at high energies. The joint $1e1p$ and $1\mu1p$ covariance and correlation matrices are shown in figure 5-8. The large off-block-diagonal correlations between the $1e1p$ and $1\mu1p$ samples will be exploited in section 5.1.3.




Figure 5-7: The uncertainty in each bin of the E_ν^{range} distribution of the $1e1p$ (figure 5-7a) and $1\mu1p$ (figure 5-7b) samples.

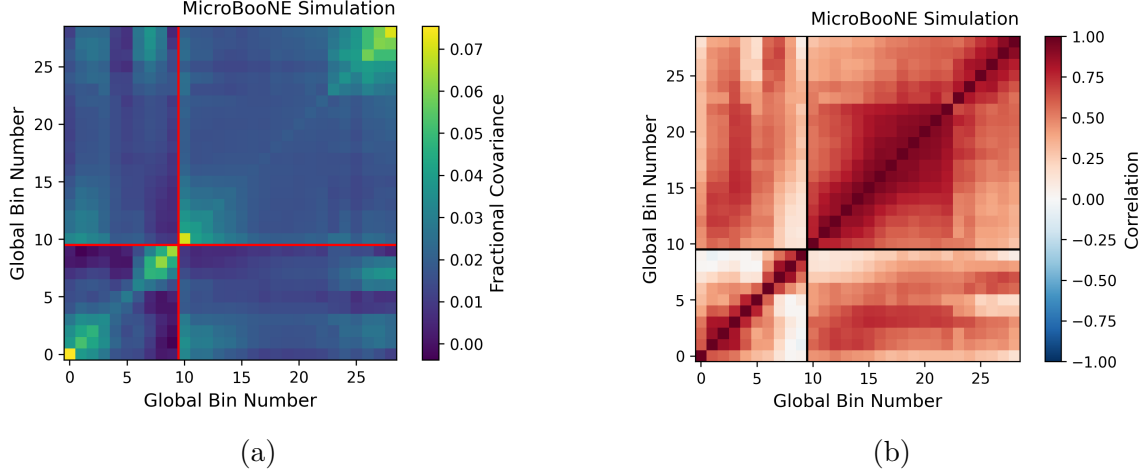


Figure 5-8: The joint covariance (figure 5-8a) and correlation (figure 5-8b) matrices for the E_ν^{range} distribution of the $1e1p$ and $1\mu1p$ samples.

5.1.3 Constraint from the $1\mu1p$ Sample

As alluded to thus far, we use the observation in the $1\mu1p$ control sample to constrain the prediction and uncertainties in the $1e1p$ signal sample. We rely on the conditional covariance method, which has also been used by the MiniBooNE collaboration [76]. We perform the constraint in the E_ν^{range} distribution of each sample. Let us represent the joint covariance matrix between both samples shown in figure 5-8a as

$$M = \begin{pmatrix} M^{ee} & M^{e\mu} \\ M^{\mu e} & M^{\mu\mu} \end{pmatrix}. \quad (5.3)$$

Given a measurement d_μ and prediction μ_μ in the $1\mu1p$ channel, the constrained prediction $\mu_{e,\text{const}}$ and covariance matrix $M^{ee,\text{const}}$ are given by [170]

$$\begin{aligned} \mu_{e,\text{const}} &= \mu_e + M^{e\mu} \cdot (M^{\mu\mu})^{-1} \cdot (d_\mu - \mu_\mu), \\ M^{ee,\text{const}} &= M^{ee} - M^{e\mu} \cdot (M^{\mu\mu})^{-1} \cdot M^{\mu e}, \end{aligned} \quad (5.4)$$

where μ_e is the nominal prediction in the $1e1p$ channel before the constraint.

Figure 5-9 shows the comparison between data and prediction in the $1\mu1p$ E_ν^{range} distribution. This observation is used as input to the constraint procedure in equation (5.4). The diagonal uncertainties in the $1e1p$ E_ν^{range} distribution are reduced by

up to a factor of ~ 2 across most bins, as shown in figure 5-10. In figure 5-11 we compare observed data to the final constrained prediction and uncertainties in the $1e1p$ E_ν^{range} distribution. As can be seen, the constraint procedure slightly increases the $1e1p$ prediction at lower E_ν^{range} values, though this is not a large effect. The result shown in figure 5-11 forms the basis of the statistical tests discussed in section 5.2.




Figure 5-9: The E_ν^{range} distribution in the $1\mu 1p$ channel, comparing data to the MC prediction.




Figure 5-10: Fractional systematic uncertainty in the $1e1p$ E_ν^{range} distribution before and after the $1\mu 1p$ constraint.




Figure 5-11: Comparison between data and prediction in the $1e1p$ E_ν^{range} distribution after applying the $1\mu1p$ constraint procedure.

5.1.4 Blinded Analysis Approach

Here we briefly cover the blind analysis procedure and related cross-checks performed before examining data in the $1e1p$ signal sample. The main strategy here was to use a sequential unblinding approach. Data were first examined for events with a low BDT score (< 0.7) or a high reconstructed neutrino energy (> 700 MeV). Next, we examined data with a medium BDT score (within $[0.7, 0.95]$) or medium energy ($500 - 700$ MeV). Only after establishing good agreement between data and prediction in all of these samples were potential signal events first examined, where signal events are defined to have a $1e1p$ BDT score greater than 0.95 and reconstructed neutrino energy below 500 MeV. This sequence was chosen because sensitivity to the eLEE model is minimized at lower BDT scores and higher neutrino energies. Moving on to each subsequent stage required signoff from the entire MicroBooNE collaboration. Figure 5-12 shows the comparison between data and prediction in the $1e1p$ BDT ensemble average score distribution for this “medium BDT score” sample. Figure 5-6 is another example of one of the “medium BDT score” plots.

We also explored a sample that made a number of cuts on kinematic variables in order to mimic the $1e1p$ BDT ensemble cut. As humans are poorer optimizers than BDTs, this sample contained a much larger rate of ν_μ background events compared

to the official signal selection. Good agreement between data and prediction was observed across all variables explored within this sample, thus giving confidence in the MC description of backgrounds to the $1e1p$ signal sample.

Finally, before any ν_e analysis was able to look at data in their signal sample, we each had to pass a series of fake data tests. In these tests, five different data-like MC samples were generated with various unknown alterations to the nominal MC. These included, for example, the presence of a large ν_e eLEE signal and alterations to underlying cross section models. All three ν_e analyses ran through the statistical results presented in section 5.2 for each fake data sample. In all cases, the statistical conclusions determined from each fake dataset were consistent with the truth-level injected events. At this point, the ν_e analyses (including the two-body CCQE analysis) were ready to move on to the real signal dataset.




Figure 5-12: Comparison between data and prediction in the $1e1p$ BDT ensemble average score distribution within the range $[0.7, 0.95]$.

5.2 Statistical Interpretation

We perform three different statistical tests of the eLEE model: a χ^2 goodness-of-fit test, a $\Delta\chi^2$ two-hypothesis test, and a signal strength scaling test. We use the

combined Neyman-Pearson (CNP) chi-square test statistic [210],

$$\chi_{\text{CNP}}^2 = (\vec{d} - \vec{\mu})^T (M_{\text{CNP}}^{\text{stat}} + M^{\text{sys}})^{-1} (\vec{d} - \vec{\mu}), \quad (5.5)$$

where \vec{d} and $\vec{\mu}$ are the observation and prediction for a particular distribution, M^{sys} is the systematic covariance matrix described in section 5.1.2, and

$$(M_{\text{CNP}}^{\text{stat}})_{ij} \equiv \frac{3\delta_{ij}}{\frac{1}{d_i} + \frac{2}{\mu_i}}. \quad (5.6)$$

This formulation of the statistical error is meant to more closely approximate the Poisson likelihood, which is important given the low-statistics nature of the *1e1p* analysis. Unlike the Poisson likelihood, the CNP chi-square test statistic is compatible with the covariance matrix formalism developed for our systematic uncertainty treatment.

All statistical tests use the constrained *1e1p* prediction and uncertainties shown in figure 5-11 when comparing to the 25 observed data events. We define two distinct hypotheses which we will refer to in the following discussion: H_0 , the prediction without the eLEE model, considering only intrinsic ν_e events and non- ν_e backgrounds (i.e. the stacked histogram in figure 5-11); and H_1 , the prediction including the eLEE model (i.e. the dashed blue line in figure 5-11). Throughout this section, we report probabilities, or p -values, calculated using the following frequentist method. To test the distribution of a specific variable under a given hypothesis, 10^5 pseudo-experiments are generated by sampling the multivariate distribution defined by the prediction and covariance matrix of the hypothesis. This gives a predicted rate in each bin of the distribution; we use this to sample from a Poisson distribution to get an integer number of “observed” events in each bin. We calculate the desired test statistic for each pseudo-experiment to build up the probability distribution of that test statistic, which incorporates both systematic and statistical uncertainties. Then, we can compare the observed value of the test statistic in the real dataset to the calculated probability distribution to determine the probability of observing a more

extreme value–this defines the p -value.

5.2.1 Goodness of Fit

For the first test, we evaluate the goodness of fit in the $1e1p$ sample for H_0 and H_1 using the χ_{CNP}^2 test statistic. This is done both before and after the $1\mu1p$ constraint. We also consider the goodness-of-fit in both the full analysis range, $200 < E_\nu^{\text{range}} [\text{MeV}] < 1200$, and the low energy region $200 < E_\nu^{\text{range}} [\text{MeV}] < 500$, where the eLEE model prediction is highest. Table 5.2 shows the χ_{CNP}^2 per degree-of-freedom (DOF) for each of these scenarios, as well as the corresponding p -values calculated via the frequentist method described above.

The first thing to note there is significant tension with H_1 , especially after the $1\mu1p$ constraint, as indicated by the low p -values. This is true both over the full analysis energy range and in the low energy range. After the constraint, there is greater than 3σ tension between H_1 and the observed data across the full energy range.

There is also disagreement between H_0 and the observed data, though not to the same level as the H_1 case. Over the full analysis energy range, the p -values in table 5.2 indicate 2.1σ (2.4σ) tension between H_0 and the data before (after) the $1\mu1p$ constraint. This seems to come from two sources: the deficit in data compared to prediction below 500 MeV in figure 5-11, and the excess in the 800-900 MeV bin in figure 5-11. The first of these is possibly a real effect–similar low-energy deficits are observed by all three ν_e analyses [116]. It is also observed in the $E_\nu^{QE-\ell}$ distribution in figure 5-4d and is therefore not likely to be related to proton misreconstruction. Our Michel electron sample covered in section 4.4.3 indicates robust modeling of low-energy electron showers, so this is also unlikely to be the culprit. The excess in the 800-900 MeV bin, however, appears to be a statistical fluctuation. The other variables examined in this analysis all show good agreement between data and MC and do not have similar narrow excesses. For instance, the E_e , E_p , θ_e , and $E_\nu^{QE-\ell}$ distributions in figure 5-4 have a $\chi_{\text{CNP}}^2/\text{DOF}$ (p -value) of 10.68/10 (0.42), 5.77/10 (0.80), 10.19/10 (0.48), and 11.46/10 (0.52), respectively. The $\chi_{\text{CNP}}^2/\text{DOF}$ and p -

value of all 36 variables examined in this analysis are reported in the Supplemental Material of Ref. [148].

Nominal Predictions				
Range	H_0		H_1	
	$\chi^2_{\text{CNP}}/\text{DOF}$	p -value	$\chi^2_{\text{CNP}}/\text{DOF}$	p -value
200–500 MeV	6.06/3	0.138	8.30/3	0.053
200–1200 MeV	23.02/10	0.024	25.37/10	0.014
Constrained Predictions				
Range	H_0		H_1	
	$\chi^2_{\text{CNP}}/\text{DOF}$	p -value	$\chi^2_{\text{CNP}}/\text{DOF}$	p -value
200–500 MeV	7.91/3	0.075	17.3/3	0.002
200–1200 MeV	25.28/10	0.014	36.35/10	5.0×10^{-4}

Table 5.2: Results from goodness-of-fit tests comparing observed $1e1p$ data to the H_0 and H_1 predictions, reported via the χ^2_{CNP} test statistic and the frequentist p -value. The top half of the table considers the nominal prediction and uncertainties before the $1\mu1p$ constraint described in section 5.1.3, while the bottom half considers the post-constraint prediction and uncertainties.

5.2.2 Two Hypothesis Test

The next test we perform is a two-hypothesis $\Delta\chi^2$ test between H_0 and H_1 . Specifically, we use the test statistic

$$\Delta\chi^2 = \chi^2_{\text{CNP},H_0} - \chi^2_{\text{CNP},H_1}. \quad (5.7)$$

Thus, a positive (negative) value indicates better agreement with H_1 (H_0). For this test, we consider the post- $1\mu1p$ -constraint prediction and errors for the $1e1p$ sample. Figure 5-13 shows the distribution of this test statistic under both hypotheses, generated using the frequentist method described above. As expected, the H_0 distribution peaks below zero while the H_1 distribution peaks above zero. These distributions can be used to determine the one-sided exclusion sensitivity of each hypothesis in the two-body CCQE analysis. If we denote the median $\Delta\chi^2$ value under H_1 as $(\Delta\chi^2)_{H_1}^{\text{med}}$, the probability of H_0 generating a larger $\Delta\chi^2$ than $(\Delta\chi^2)_{H_1}^{\text{med}}$ is 0.003. This corresponds to an H_0 exclusion sensitivity of 2.7σ . Conversely, if we define $(\Delta\chi^2)_{H_0}^{\text{med}}$ in

an analogous way, the probability of H_1 generating a *smaller* $\Delta\chi^2$ than $(\Delta\chi^2)_{H_0}^{\text{med}}$ is 0.017, corresponding to an H_1 exclusion sensitivity of 2.1σ .

As shown in figure 5-13, the observed data result in a test statistic of $\Delta\chi^2 = -11.08$. The probability of observing a smaller $\Delta\chi^2$ under H_0 and H_1 is 0.020 and 1.6×10^{-4} , respectively, corresponding to a Gaussian significance of 2.1σ and 3.6σ . Thus, the data exhibits mild tension with H_0 and significant tension with H_1 . A similar conclusion was reached for the goodness-of-fit test; the tension with H_0 likely comes from the same sources discussed in section 5.2.1. Due to the fact that non-negligible tension is observed with respect to both hypotheses, we report our exclusion significance for H_1 using the CL_s method [211, 212]. The CL_s test statistic is defined as

$$\text{CL}_s = \frac{p_{H_1}}{p_{H_0}}, \quad (5.8)$$

where $\text{CL}_s = 0.008$ for our observed data. Using this method, we reject H_1 in favor of H_0 with a Gaussian significance of 2.4σ .




Figure 5-13: Distributions of the $\Delta\chi^2$ test statistic defined in equation (5.7) for H_0 (red) and H_1 (blue), calculated by generating 10^5 pseudo-experiments under each hypothesis. The $\Delta\chi^2$ value of the data is also shown.

5.2.3 Signal Strength Scaling Test

The final statistical test we perform on our $1e1p$ dataset is a one-parameter signal strength scaling test. Specifically, we multiply by the eLEE weights in figure 4-1b by an overall normalization factor x_{LEE} which can vary within $[0, \infty]$. In this construction, $x_{\text{LEE}} = 0$ corresponds to H_0 and $x_{\text{LEE}} = 1$ corresponds to H_1 . We can use the Feldman-Cousins procedure [25] along with the test statistic

$$\Delta\chi^2(x_{\text{LEE}}) = \chi_{\text{CNP}}^2(x_{\text{LEE}}) - \min_{x_{\text{LEE}} \in [0, \infty]} \{\chi_{\text{CNP}}^2(x_{\text{LEE}})\} \quad (5.9)$$

to determine confidence intervals on x_{LEE} . We can use the Asimov dataset, in which the observation is exactly the predicted value in each bin, to determine the sensitivity to each hypothesis in this test. Under H_0 , the expected upper bound on x_{LEE} at the 90% (2σ) confidence level is 0.75 (0.98). Under H_1 , the expected confidence interval at the 1σ (2σ) confidence level is $[0.53, 1.66]$ ($[0.28, 2.67]$), and the expected sensitivity to rule out $x_{\text{LEE}} = 0$ is 2.8σ .

Figure 5-14 shows the confidence intervals on x_{LEE} derived from the observed data. The best-fit value is $x_{\text{LEE}} = 0$, which is expected given the deficit in data at the lowest neutrino energies where the eLEE model weights are the largest. The 90% (2σ) upper bound on x_{LEE} is 0.25 (0.38). These upper bounds are stronger than the expected confidence intervals under H_0 , which is mainly due to the aforementioned deficit at low neutrino energies.

5.3 Discussion and Outlook

The observation in the $1e1p$ signal channel of the two-body CCQE analysis is inconsistent with a MiniBooNE-like excess of ν_e CCQE interactions. Under a two-hypothesis test, we rule out the nominal eLEE model (H_1) at the 2.4σ using the CL_s method. We can set an upper bound on the signal strength scaling parameter x_{LEE} of 0.38 at the 2σ level.

This is consistent with the results from all three MicroBooNE ν_e analyses [116].




Figure 5-14: Confidence intervals on x_{LEE} calculating using the Feldman-Cousins procedure. The solid and dotted lines indicate the confidence level with which a given x_{LEE} is disfavored, calculated using the Feldman-Cousins method [25] and Wilks theorem [26], respectively. The MiniBooNE statistical and systematic errors are shown as a band around $x_{\text{LEE}} = 1$.

As shown in figure 5-15a, three of the four signal channels from the three ν_e analyses observe a deficit of selected events at the lowest neutrino energies, where the eLEE model is supposed to peak. The exception is the $1e0p0\pi$ channel from the MiniBooNE-like sample; however, this is the least sensitive of the four channels, as it is dominated by non- ν_e backgrounds. Figure 5-15b shows the results from the signal strength scaling test in all four channels. With the exception of the $1e0p0\pi$ channel, all analyses rule out $x_{\text{LEE}} = 1$ at greater than the 2σ confidence level. The most sensitive analyses observe an upper bound $x_{\text{LEE}} \lesssim 0.5$ at the 2σ confidence level.

The first MicroBooNE results also included a search for $\Delta \rightarrow N\gamma$ decays as an explanation for the MiniBooNE excess. The analysis did not observe an excess of $\Delta \rightarrow N\gamma$ candidate events and ruled out the required enhancement to explain the MiniBooNE anomaly at the 95% confidence level. The full details of this analysis can be found in Ref. [80].

Thus, the MicroBooNE results disfavor both an enhancement of ν_e CC interac-

tions and $\Delta \rightarrow N\gamma$ decays as the sole explanation of the MiniBooNE excess. Potential explanations of the MiniBooNE anomaly must now remain consistent with these MicroBooNE results. One promising avenue involves dark sector models which produce a photon or e^+e^- pair in the MiniBooNE detector. Such event signatures have not yet been explored by MicroBooNE, though this work is in progress. Crefch:neutrissimos explores a photon-based explanation of MiniBooNE, involving the decay of a heavy neutral lepton to a single photon via a transition magnetic moment. The theoretical and experimental neutrino physics communities are actively working to develop innovative models to explain the MiniBooNE LEE and equally innovative measurements to test those models.

The following subsections close with two projects I co-lead with various collaborators to follow up on the state of the MiniBooNE excess after the MicroBooNE results. The first examines the $3+1$ model simultaneously in MiniBooNE and MicroBooNE [102]. The second explores the implications of a $\bar{\nu}_e$ -based MiniBooNE LEE for the MicroBooNE results [201].

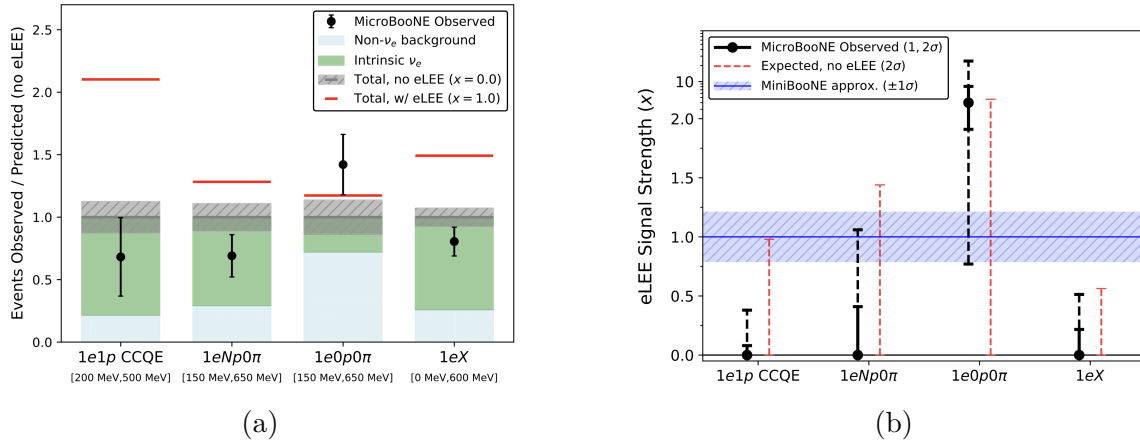


Figure 5-15: Figure 5-15a shows the observation compared to the nominal (H_0) prediction in all four signal channels from the three MicroBooNE ν_e analyses, including statistical errors on the data points and systematic errors on the prediction. The eLEE prediction (H_1) is also indicated by the red line. Figure 5-15b shows the observed 1σ and 2σ confidence intervals on x_{LEE} from all four signal channels. The 2σ expected sensitivity of each channel is shown in red.

5.3.1 Publication: *MiniBooNE and MicroBooNE Combined Fit to a 3 + 1 Sterile Neutrino Scenario*

The initial results from the MicroBooNE ν_e analyses did not explicitly test short-baseline $\nu_\mu \rightarrow \nu_e$ oscillations from the 3 + 1 model. Shortly after the MicroBooNE results came out, two studies appeared assessing the sensitivity of the MicroBooNE ν_e analyses to the 3 + 1 model [188, 213]. Following up on this, the MiniBooNE collaboration released the first combined fit to a 3 + 1 model using MicroBooNE and MiniBooNE data [102]. This study was led by myself and Austin Schneider. The full *Physical Review Letters* publication is included below. We considered the two-body CCQE and inclusive MicroBooNE ν_e analyses, as these placed the most stringent constraints on the eLEE model and provided data releases with the requisite detail to perform a 3 + 1 fit [116]. The most important result from this study is Figure 2, which indicates that there are still allowed regions in 3 + 1 parameter space at the 3σ confidence level in the MiniBooNE-MicroBooNE combined fit. This was also the first MiniBooNE analysis to account for all three oscillation probabilities in equation (1.36): ν_e appearance, ν_μ disappearance, and ν_e disappearance. Previous 3 + 1 fits to MiniBooNE data have only considered ν_e appearance, as this is the dominant effect in the electron-like channel. However, ν_e and ν_μ disappearance can have a non-negligible impact on the intrinsic ν_e and misidentified ν_μ background rate in the MiniBooNE electron-like prediction. The ν_e disappearance probability is also important when considering MicroBooNE data, as the prediction in each MicroBooNE ν_e analysis is dominated by intrinsic ν_e interactions [148, 170]. MicroBooNE has since released official constraints on the 3 + 1 model using data from the inclusive ν_e analysis [214]—these constraints are consistent with our implementation.

MiniBooNE and MicroBooNE Combined Fit to a 3 + 1 Sterile Neutrino Scenario

A. A. Aguilar-Arevalo,¹⁴ B. C. Brown,⁵ J. M. Conrad,¹³ R. Dharmapalan,^{1,7} A. Diaz,¹³ Z. Djurcic,² D. A. Finley,⁵ R. Ford,⁵ G. T. Garvey,¹⁰ S. Gollapinni,¹⁰ A. Hourlier,¹³ E.-C. Huang,¹⁰ N. W. Kamp,¹³ G. Karagiorgi,⁴ T. Katori,¹² T. Kobilarcik,⁵ K. Lin,^{4,10} W. C. Louis,¹⁰ C. Mariani,¹⁶ W. Marsh,⁵ G. B. Mills,^{10,†} J. Mirabal-Martinez,¹⁰ C. D. Moore,⁵ R. H. Nelson,^{3,*} J. Nowak,⁹ Z. Pavlovic,⁵ H. Ray,⁶ B. P. Roe,¹⁵ A. D. Russell,⁵ A. Schneider[Ⓞ],¹³ M. H. Shaevitz,⁴ J. Spitz,¹⁵ I. Stancu,¹ R. Tayloe,⁸ R. T. Thornton,¹⁰ M. Tzanov,^{3,11} R. G. Van de Water,¹⁰ D. H. White,^{10,†} and E. D. Zimmerman³

(The MiniBooNE Collaboration)¹

¹University of Alabama, Tuscaloosa, Alabama 35487, USA

²Argonne National Laboratory, Argonne, Illinois 60439, USA

³University of Colorado, Boulder, Colorado 80309, USA

⁴Columbia University, New York, New York 10027, USA

⁵Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA

⁶University of Florida, Gainesville, Florida 32611, USA

⁷University of Hawaii, Manoa, Honolulu, Hawaii 96822, USA

⁸Indiana University, Bloomington, Indiana 47405, USA

⁹Lancaster University, Lancaster LA1 4YB, United Kingdom

¹⁰Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA

¹¹Louisiana State University, Baton Rouge, Louisiana 70803, USA

¹²King's College London, London WC2R 2LS, United Kingdom

¹³Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

¹⁴Instituto de Ciencias Nucleares, Universidad Nacional Autónoma de México, CDMX 04510, México

¹⁵University of Michigan, Ann Arbor, Michigan 48109, USA

¹⁶Center for Neutrino Physics, Virginia Tech, Blacksburg, Virginia 24061, USA



(Received 5 January 2022; revised 9 September 2022; accepted 28 September 2022; published 8 November 2022)

This Letter presents the results from the MiniBooNE experiment within a full “3 + 1” scenario where one sterile neutrino is introduced to the three-active-neutrino picture. In addition to electron-neutrino appearance at short baselines, this scenario also allows for disappearance of the muon-neutrino and electron-neutrino fluxes in the Booster Neutrino Beam, which is shared by the MicroBooNE experiment. We present the 3 + 1 fit to the MiniBooNE electron-(anti)neutrino and muon-(anti)neutrino data alone and in combination with MicroBooNE electron-neutrino data. The best-fit parameters of the combined fit with the exclusive charged-current quasielastic analysis (inclusive analysis) are $\Delta m^2 = 0.209 \text{ eV}^2(0.033 \text{ eV}^2)$, $|U_{e4}|^2 = 0.016(0.500)$, $|U_{\mu 4}|^2 = 0.500(0.500)$, and $\sin^2(2\theta_{\mu e}) = 0.0316(1.0)$. Comparing the no-oscillation scenario to the 3 + 1 model, the data prefer the 3 + 1 model with a $\Delta\chi^2/\text{d.o.f.} = 24.7/3(17.3/3)$, a $4.3\sigma(3.4\sigma)$ preference assuming the asymptotic approximation given by Wilks’s theorem.

DOI: [10.1103/PhysRevLett.129.201801](https://doi.org/10.1103/PhysRevLett.129.201801)

Introduction.—The MiniBooNE low-energy excess (LEE) is a long-standing anomaly in neutrino physics. This excess of electronlike events was observed in the muon-neutrino dominated flux from the Booster Neutrino Beam (BNB), and is most significant between 200 and 600 MeV in reconstructed neutrino energy. Initially

reported in 2007 [1], the excess reached a significance of 4.8σ in the energy range $200 < E_{\nu}^{\text{QE}} < 1250 \text{ MeV}$ with the full MiniBooNE ν and $\bar{\nu}$ dataset [2]. We note that this significance is derived from a direct comparison between MiniBooNE data and the Standard Model (SM) prediction and is thus independent of any physics model, including the 3 + 1 model explored in this Letter. A wide range of explanations for the excess have been put forward, but the initial, and still most-referenced, new physics explanations invoke $\nu_{\mu} \rightarrow \nu_e$ oscillations.

The BNB flux is produced through 8 GeV protons impinging on a beryllium target that is located inside a

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.

magnetic focusing horn, which can reverse polarity to run in neutrino or antineutrino mode, followed by a 50 m meson decay pipe. The MiniBooNE detector, which is a 450 ton fiducial mass, mineral-oil-based Cherenkov detector, is located 541 m downstream of the beryllium target. The detector is sensitive to neutrinos with energies between 100 MeV and 3 GeV. This combination of energy and baseline makes MiniBooNE an ideal experiment to probe the appearance of electron neutrinos from $\nu_\mu \rightarrow \nu_e$ oscillations in a mass-squared splitting region Δm^2 greater than $1 \times 10^{-2} \text{ eV}^2$. The full dataset taken in a series of runs between 2002 and 2019 yields a 1σ allowed region in Δm^2 between 0.04 and 0.4 eV^2 , with mixing angles varying from 1.0 to 0.01 [2].

These mass-squared splittings are more than an order of magnitude larger than the splitting of atmospheric neutrino oscillations, $\Delta m_{\text{atmos}}^2 \approx 2.5 \times 10^{-3} \text{ eV}^2$ [3]—associated with the largest mass splitting in three neutrino oscillation models. Therefore, to accommodate such oscillations, it is necessary to postulate the existence of a fourth neutrino mass and a fourth neutrino flavor that must be non-weakly-interacting (or “sterile”) to avoid constraints from Z decay [4]. In such a model, the sterile neutrino flavor and the three active flavors are connected to a fourth mass state through an extension of the Pontecorvo–Maki–Nakagawa–Sakata mixing matrix. Such a model introduces a combination of three possible experimental signatures: (1) electron flavor disappearance to other flavors, leading to fewer ν_e events than expected (“ $\nu_e \rightarrow \nu_e$ ”), (2) muon flavor disappearance to other flavors (“ $\nu_\mu \rightarrow \nu_\mu$ ”) reducing the ν_μ rate, and (3) $\nu_\mu \rightarrow \nu_e$ appearance, where an excess of ν_e events would be observed. Past MiniBooNE $\nu_\mu \rightarrow \nu_e$ appearance analyses have assumed that the ν_e and ν_μ disappearance effects were negligible. However, for the mass-squared splitting and mixing angles we are concerned with, ν_e disappearance can reduce the intrinsic ν_e background contribution by up to 80%, and ν_μ disappearance can decrease the ν_μ event rate by up to 80%. Neglecting these effects has been considered to be an overly simplified approach. External analyzers have investigated the difference between this approach and an analysis with a full treatment of the $3 + 1$ model [5–7] and have explored the effects other nuclear models on these results [5]. In response to this, in this Letter we expand the analyses of the full MiniBooNE datasets and simulation samples to present the first full $3 + 1$ sterile neutrino oscillation model by the collaboration.

In 2015, the MicroBooNE experiment joined the MiniBooNE experiment as a user of the BNB beamline. The MicroBooNE experiment was designed with the primary goal of investigating the LEE by using the detailed information from its liquid-argon time-projection chamber (LArTPC) to distinguish between electron induced events and photon induced events. This allows the rejection of many misidentified backgrounds in the MiniBooNE dataset. MicroBooNE has recently released results of a search

for a generic ν_e excess, assuming the median shape of the MiniBooNE excess, in a strategy that is agnostic to particular oscillation models. External analyses have applied more focused studies, placing limits on ν_e disappearance [8], expanding the MicroBooNE analysis to all systematically allowed shapes of the MiniBooNE excess [9], and considering how the MicroBooNE data constrain the parameters of a $3 + 1$ sterile neutrino model [9]. However, until now, there has been no MiniBooNE–MicroBooNE combined analysis.

Because MiniBooNE and MicroBooNE share the same beamline, we can use MiniBooNE tools to perform a combined fit to the data of the two experiments. On the other hand, because the detectors are substantially different, the two experiments have complementary capabilities. MicroBooNE is an 85 ton active mass LArTPC [10], which allows for detailed reconstruction of neutrino interactions that is not possible using the MiniBooNE Cherenkov detector.

The MiniBooNE experiment has a large sample size, but relatively high backgrounds from misidentification backgrounds that dominate MiniBooNE’s electron-neutrino sample. The MicroBooNE experiment uses a relatively small detector, but can remove most misidentification backgrounds [11]. The imaging capability of the LArTPC has allowed the MicroBooNE experiment to select three ν_e charged-current (CC) samples [12]: a high purity exclusive sample of charged-current quasielastic (CCQE) interactions [11], a semi-inclusive sample of pionless interactions [13], and a fully inclusive sample [14]. The MicroBooNE data allow for a clean test of the hypothesis that the MiniBooNE excess events are due to ν_e CC interactions. We consider only the CCQE and inclusive analyses, as the former has the lowest systematic uncertainty and the latter has the largest sample size. The CCQE analysis uses deep-learning-based reconstruction, while the inclusive analysis uses wire-cell-based reconstruction; thus they are hereafter identified in figures and tables by the shorthand “DL” and “WC”, respectively. In this Letter, we present the first MiniBooNE–MicroBooNE combined fits to a $3 + 1$ model.

Fit details.—The model of interest is a three-active plus one-sterile neutrino model called $3 + 1$. This model expands the 3×3 neutrino mixing matrix to 4×4 ,

$$U_{3+1} = \begin{bmatrix} U_{e1} & U_{e2} & U_{e3} & U_{e4} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} & U_{\mu4} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} & U_{\tau4} \\ U_{s1} & U_{s2} & U_{s3} & U_{s4} \end{bmatrix}. \quad (1)$$

In such a model, both ν_μ and ν_e disappearance is expected to occur with the same Δm^2 as the $\nu_\mu \rightarrow \nu_e$ appearance signal, as long as both U_{e4} and $U_{\mu4}$ are nonzero. The three processes are related through their effective mixing angles, which are expressed as

$$\begin{aligned}
 \sin^2(2\theta_{\mu\mu}) &= 4(1 - |U_{\mu 4}|^2)|U_{\mu 4}|^2, \\
 \sin^2(2\theta_{ee}) &= 4(1 - |U_{e 4}|^2)|U_{e 4}|^2, \\
 \sin^2(2\theta_{e\mu}) &= 4|U_{e 4}|^2|U_{\mu 4}|^2,
 \end{aligned} \tag{2}$$

which appear within the oscillation probability formulas,

$$\begin{aligned}
 P(\nu_\mu \rightarrow \nu_e) &= \sin^2 2\theta_{\mu e} \sin^2(\Delta m_{41}^2 L/E), \\
 P(\nu_e \rightarrow \nu_e) &= 1 - \sin^2 2\theta_{ee} \sin^2(\Delta m_{41}^2 L/E), \\
 P(\nu_\mu \rightarrow \nu_\mu) &= 1 - \sin^2 2\theta_{\mu\mu} \sin^2(\Delta m_{41}^2 L/E).
 \end{aligned} \tag{3}$$

There are three physics parameters in the $3 + 1$ model relevant to these two experiments: the sterile mass splitting $\Delta m_{4i}^2 \equiv \Delta m^2$ (where we assume degeneracy for $i \in \{1, 2, 3\}$) and the two mixings of the new mass eigenstate to the electron weak eigenstate $|U_{e4}|^2$ and muon weak eigenstate $|U_{\mu 4}|^2$. Different combinations of these parameters will induce different rates of ν_e appearance, as well as ν_μ and ν_e disappearance in the MiniBooNE and MicroBooNE detectors. In each case, the oscillation probability depends upon the true neutrino energy E and baseline of each event L .

The oscillation prediction in MiniBooNE is determined by a simple reweighting of the MiniBooNE $\nu_\mu \rightarrow \nu_e$ simulation using the oscillation formulas [Eqs. (3)]. This direct method is not possible for the MicroBooNE analyses, as only limited simulation information for each analysis is available [15,16]. Instead, for MicroBooNE, we use the MiniBooNE BNB simulation to obtain a ratio between the nominal intrinsic ν_e background prediction and the ν_e appearance prediction at the MicroBooNE baseline as a function of true neutrino energy, using the BNB flux prediction at the MicroBooNE location. This ratio, combined with the intrinsic ν_e simulation provided by MicroBooNE allows us to obtain a ν_e appearance prediction in MicroBooNE. We use the same procedure to account for ν_e disappearance in both analyses and ν_μ disappearance in the inclusive analysis. We neglect ν_μ disappearance in the MicroBooNE $1e1p$ CCQE prediction, as the ν_μ background contamination in MicroBooNE's $1e1p$ analysis is subdominant and the simulation information for the ν_μ contribution is not provided by MicroBooNE. In the MicroBooNE inclusive analysis, we consider only four of the seven channels: the ν_e and ν_μ CC fully contained (FC) and partially contained (PC) samples. This is because energy reconstruction information is not provided for the remaining three π^0 -based samples. We also note that $\nu_\mu \rightarrow \nu_\tau$ neutral-current backgrounds in MiniBooNE's electron-neutrino measurement are not included in the prediction; however, this effect is expected to be small. An example of this oscillation prediction is shown in Fig. 1.

For the MiniBooNE likelihood, we compare the fixed observation to the theoretical expectation with a multivariate normal distribution that includes systematic uncertainties, Poisson statistical uncertainties on the expectation, and finite Monte Carlo statistical uncertainties. With the large MiniBooNE sample size, the multivariate normal distribution is a reasonable approximation for the likelihood. The MiniBooNE systematic errors of this analysis remain the same as in [2], with one exception. The correlated systematic errors from uncertainties in the MiniBooNE optical model are limited to the three principal components of the corresponding covariance matrix with the largest eigenvalues, and the remaining optical model errors are assumed to be uncorrelated with no covariance among energy bins. For each MicroBooNE analysis, we use a Poisson-derived likelihood that accounts for finite Monte Carlo size [17]; additionally, the expectation in each bin is treated as a nuisance parameter that is constrained by the systematics covariance matrix [15,16]. The total likelihood is then composed of these two experimental likelihoods; for more details see S1 of the Supplemental Material. We note that, although the same beamline simulation is used to derive systematic uncertainties for both experiments, because of technical limitations the fit presented here accounts for these uncertainties as if they were uncorrelated between the two experiments. The inclusion of information from the ν_μ samples of both MiniBooNE and MicroBooNE indirectly constrains the ν_e predictions of the two experiments in a correlated manner. This is handled directly for the inclusive analysis, as the ν_μ CC FC-PC samples are included in the fit (accounting for ν_μ appearance). For the CCQE analysis, we allow the MicroBooNE $\nu_\mu 1\mu 1p$ measurement to constrain the MicroBooNE $\nu_e 1e 1p$ prediction and uncertainties, and do not account for oscillations in MicroBooNE's $\nu_\mu 1\mu 1p$ prediction. Ignoring ν_μ disappearance in MicroBooNE's CCQE analysis is a reasonable assumption for small $U_{\mu 4}$ given the limited sample size from this analysis.

Results.—With the methods described in the preceding section, we can examine the MiniBooNE LEE in the context of a $3 + 1$ sterile neutrino model, both with the MiniBooNE data alone and together with the MicroBooNE electron-neutrino data. We show the no-oscillation SM prediction as a dashed purple line in Fig. 1. In the SM case, the MiniBooNE prediction lies substantially below the data in the electron-neutrino channel. For the MicroBooNE CCQE analysis, the data lie scattered above and below the SM prediction, in part due to the small sample size. For the MicroBooNE inclusive analysis, the data lie below the SM prediction across most of the energy range. The disparity between the data and SM prediction in MiniBooNE shows the inability of the SM to accommodate the MiniBooNE low-energy excess in the electron-neutrino data, while remaining in agreement with the MiniBooNE muon-neutrino data.

In contrast to the SM, the $3 + 1$ oscillation model provides the additional freedom necessary to potentially

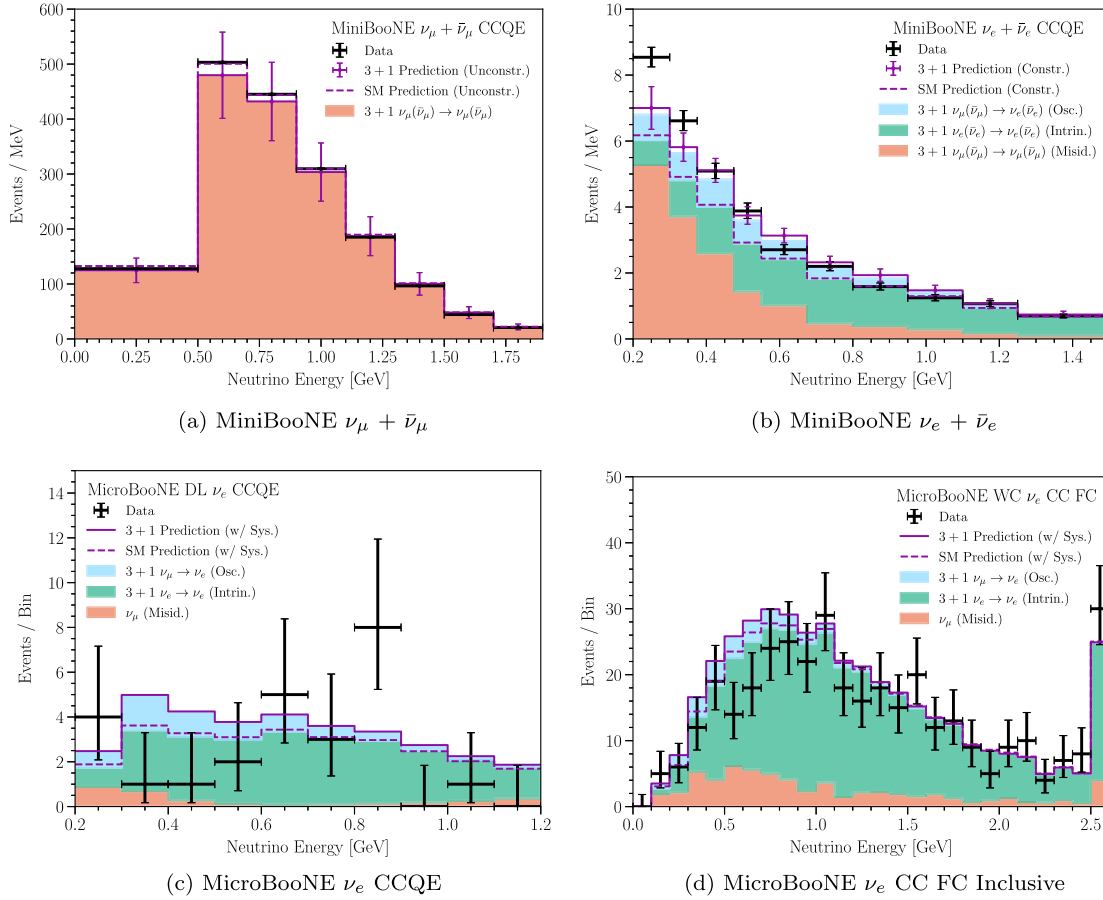


FIG. 1. Comparison between data and prediction for each experiment, showing the prediction from both the SM and the 3 + 1 model. (a)–(c) The MiniBooNE and MicroBooNE DL consider the 3 + 1 “Combination (DL)” fit parameters of Table I, while (d) the MicroBooNE WC figure considers the “Combination (WC)” fit parameters. Black crosses show the observed data and statistical error, and stacked histograms show the unconstrained prediction. The SM (3 + 1) prediction is represented as a dashed (solid) purple line. The error bars on the MiniBooNE 3 + 1 prediction represent systematic uncertainty. (a) Unconstrained predictions and errors. (b) Predictions and errors in purple after being constrained by the $\nu_\mu + \bar{\nu}_\mu$ data. (c),(d) Predictions after the allowed systematic variations have been fit to data, and thus do not have systematic error bars shown.

better accommodate the MiniBooNE muon-neutrino data and low-energy excess within systematic errors. In the 3 + 1 scenario, we expect $\nu_\mu \rightarrow \nu_e$ oscillations to increase the prediction in the electron-neutrino channels of both experiments, while ν_e disappearance will reduce the intrinsic electron-neutrino backgrounds, and ν_μ disappearance will reduce the muon-neutrino prediction as well as the contribution of misidentified events in the electron-neutrino observable channel. The prediction for the best-fit 3 + 1 scenario across both experiments is shown in Fig. 1, separated by component, experiment, and observable channel. Figure 1(a) compares the MiniBooNE unconstrained muon-neutrino and antineutrino prediction to observed data, where the crosses denote the unconstrained 3 + 1 prediction and the dashed line denotes the unconstrained SM prediction; here the 3 + 1 prediction is approximately 4% lower than the SM prediction in the bin with the largest expectation. Figure 1(b) compares the MiniBooNE electron-neutrino and antineutrino prediction

to data; the prediction and errors are shown after being constrained by the muon-neutrino data for the 3 + 1 and SM scenarios in purple, whereas the unconstrained 3 + 1 prediction is shown by the stacked histogram. The best-fit electron-neutrino 3 + 1 oscillation prediction is approximately 15% lower in the lowest energy bin than that reported for the best fit in the two-neutrino oscillation analysis [2].

While the best-fit 3 + 1 scenario is preferred to the no-oscillation scenario, it still cannot perfectly describe MiniBooNE’s low-energy excess, especially at the lowest energies. This is consistent with the recent MicroBooNE results, which indicate that the low-energy excess cannot be explained entirely by electron neutrinos [12]. This is also consistent with previous MiniBooNE studies indicating a forward-peaked angular distribution of the low-energy excess [2].

The best-fit 3 + 1 parameters and the $\Delta\chi^2$ between the SM and 3 + 1 scenarios are given in Table I.

TABLE I. Summary of results. The $\Delta\chi^2/\text{d.o.f.}$ in the last column compares the 3 + 1 model to the no-oscillation model.

3 + 1 fit	$ U_{e4} ^2$	$ U_{\mu4} ^2$	Δm^2	$\Delta\chi^2/\text{d.o.f.}$
MiniBooNE only	0.021	0.500	0.191	27.8/3
Combination (DL)	0.016	0.500	0.209	24.7/3
Combination (WC)	0.500	0.500	0.033	17.3/3

We obtain a best fit that includes substantial sterile-muon mixing, with $|U_{\mu4}|^2$ near 0.5, and moderate sterile-electron mixing, with $|U_{e4}|^2$ near 0.02, for both the MiniBooNE only and the CCQE combined fit. For these fits, the best-fit Δm^2 is near 0.2 eV^2 as well. The large sterile-muon mixing at the best-fit point is in tension with constraints on unitarity from some experiments in the neutrino sector [18–20], although a substantial number of neutrino experiments violate these constraints [21–25] as discussed in [8]. However, in this analysis, a broad region in parameter space is allowed within the estimated 1σ confidence region, as is visualized in Fig. 2, extending to regions of parameter space that are not in tension with unitarity constraints. The 1σ allowed region in Δm^2 and $\sin^2(2\theta_{\mu e})$ is similar to that reported in [2] and takes the form of a diagonal band because the MiniBooNE LEE spans a broad energy range and extends down to the 200 MeV boundary. The inclusive search combined fit obtains a best fit at maximal mixing, with a Δm^2 of 0.033 eV^2 , compatible with this diagonal band. The excess drives the allowed values of $\sin^2(2\theta_{\mu e})$, but large deviations from the best fit in $|U_{e4}|^2$ and $|U_{\mu4}|^2$ are allowed, provided the combination produces enough $\nu_\mu \rightarrow \nu_e$ appearance to describe the excess. This freedom is present in part because the systematic errors of the prediction allow large changes to the muon-neutrino channel with little penalty, which in turn provides only a weak constraint on $|U_{\mu4}|^2$ through ν_μ disappearance.

In both analyses, MicroBooNE’s electron-neutrino data do not exhibit an excess at the lower end of their energy spectrum, as MiniBooNE’s electron-neutrino data do, and MicroBooNE overall observes a lower event rate than predicted by the nominal no-oscillation model [12]. However, the data sample from MicroBooNE does not have the statistical power needed to rule out a 3 + 1 $\nu_\mu \rightarrow \nu_e$ explanation of the MiniBooNE low-energy excess. The observed event rate from MicroBooNE’s ν_e CCQE $1e1p$ analysis precludes very large $\nu_\mu \rightarrow \nu_e$ appearance at values of Δm^2 and $\sin^2(2\theta_{\mu e})$ higher than the MiniBooNE allowed region. This manifests in Fig. 2 (top) as a small shift in the allowed region to lower Δm^2 and lower $\sin^2(2\theta_{\mu e})$. In Fig. 1(c), the best-fit 3 + 1 oscillation prediction increases the expected number of events in a region where the MicroBooNE CCQE analysis observes a deficit, suggesting that the fit is primarily driven by the larger MiniBooNE

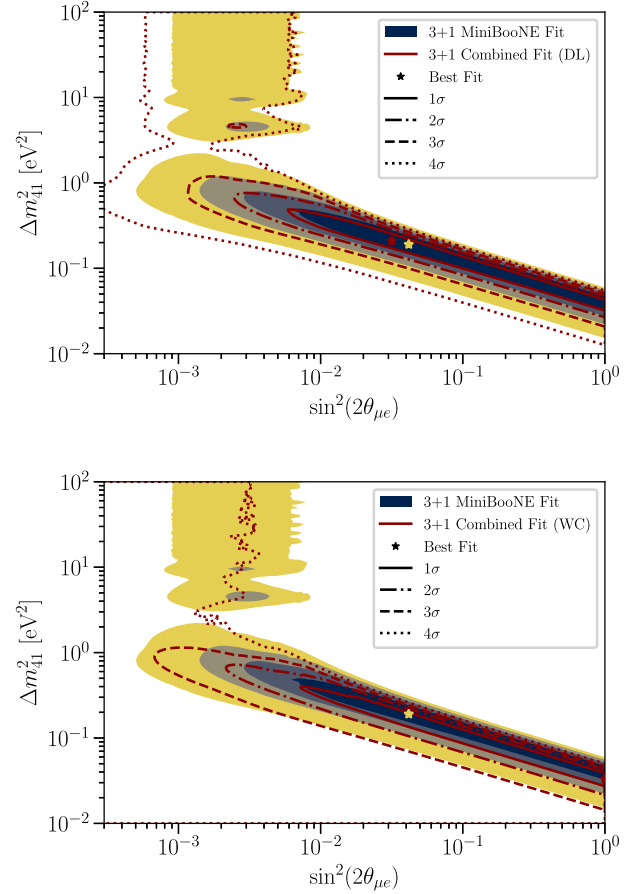


FIG. 2. The results of the MiniBooNE-only and combined fits with MicroBooNE’s CCQE sample [11] (top) and MicroBooNE’s inclusive sample [14] (bottom). The likelihood is obtained by profiling over all parameters except Δm^2 and $\sin^2(2\theta_{\mu e})$. The two best-fit points are shown as appropriately colored stars, and the contours are obtained by comparing the profile-likelihood-ratio test statistic to the asymptotic distribution provided by Wilks’s theorem, assuming a difference of 2 degrees of freedom.

data sample, in line with our expectation. This is also true for the MicroBooNE inclusive analysis, as shown in Fig. 1(d). However, the inclusive analysis provides a stronger constraint on $\nu_\mu \rightarrow \nu_e$ appearance in general. This manifests in Fig. 2 as (1) a more significant modification of the allowed regions for the combined fit and (2) a smaller $\Delta\chi^2$ between the 3 + 1 best fit and the SM in Table I, in comparison to the combined fit with the MicroBooNE CCQE analysis. More details on the combined fit with the inclusive analysis and MicroBooNE-only constraints from both analyses can be found in S2 and S3 of the Supplemental Material, respectively.

The 3 + 1 scenario is preferred over the no-oscillation model in both the MiniBooNE-only and combined-fit cases. In the MiniBooNE-only fit, we obtain a $\Delta\chi^2 = 27.8$ between the two models, whereas in the combined-fit we obtain a $\Delta\chi^2 = 24.7$ for 3 additional degrees of freedom introduced in the fit. This is smaller than the $\Delta\chi^2 = 29$

with 3 degrees of freedom reported in the two-neutrino oscillation analysis [2], representing a drop in the significance when disappearance effects are accounted for. If we assume the asymptotic approximation to the test-statistic distribution provided by Wilks's theorem [26] with a difference of 3 degrees of freedom between the models, then we obtain p values of 4.09×10^{-6} and 1.77×10^{-5} in favor of the $3 + 1$ scenario for the MiniBooNE-only and combined analyses, respectively. However, we expect the true difference in degrees of freedom between the models to be less than 3, based on both the degeneracy inherent in the $3 + 1$ model and the smaller difference in degrees of freedom observed in the two-neutrino MiniBooNE oscillation study ([2], Sec. 5). A reduction in the difference in degrees of freedom between the models would increase the significance of these two statistical tests. Therefore, we conservatively estimate that the MiniBooNE-only $3 + 1$ model test prefers the $3 + 1$ model to the SM at approximately 4.6σ , and the addition of the MicroBooNE electron-neutrino CCQE (inclusive) data reduces this significance to approximately 4.3σ (3.4σ).

Conclusion.—This Letter has explored a full $3 + 1$ sterile neutrino oscillation model within the context of results from the MiniBooNE and MicroBooNE experiments. In the MiniBooNE electronlike analysis, we consider $\nu_\mu \rightarrow \nu_e$ appearance alongside both ν_e and ν_μ disappearance. In the MicroBooNE CCQE analysis, we consider ν_e appearance and ν_e disappearance. In the MicroBooNE inclusive analysis, we consider ν_e appearance and both ν_e and ν_μ disappearance. In an analysis of the MiniBooNE-only data, we find a best fit to the $3 + 1$ model of $\Delta m^2 = 0.191 \text{ eV}^2$, $|U_{e4}|^2 = 0.021$, $|U_{\mu4}|^2 = 0.500$, and $\sin^2(2\theta_{\mu e}) = 0.0417$. A combined fit to the MiniBooNE and MicroBooNE CCQE analyses finds a best fit to the $3 + 1$ model at oscillation parameters of $\Delta m^2 = 0.209 \text{ eV}^2$, $|U_{e4}|^2 = 0.016$, $|U_{\mu4}|^2 = 0.500$, and $\sin^2(2\theta_{\mu e}) = 0.0316$. A combined fit to the MiniBooNE and MicroBooNE inclusive analyses finds a best fit to the $3 + 1$ model at oscillation parameters of $\Delta m^2 = 0.033 \text{ eV}^2$, $|U_{e4}|^2 = 0.500$, $|U_{\mu4}|^2 = 0.500$, and $\sin^2(2\theta_{\mu e}) = 1.0$. In the MiniBooNE-only analysis, the $3 + 1$ scenario is preferred over the no-oscillation case with a $\Delta\chi^2/\text{d.o.f.} = 27.8/3$, whereas in the combined analysis with MicroBooNE CCQE (inclusive) data, we obtain $\Delta\chi^2/\text{d.o.f.} = 24.7/3$ ($17.3/3$). Although the $3 + 1$ model is not a perfect description of the low-energy MiniBooNE electron-neutrino data, we find that a $3 + 1$ sterile neutrino oscillation scenario is a better description of the MiniBooNE data than the no-oscillation scenario and is not in tension with MiniBooNE's muon-neutrino data. We also find that the MicroBooNE electron-neutrino data do not rule out the allowed $3 + 1$ interpretations for the MiniBooNE data, but do reduce the significance of the result and make only a small modification to the allowed regions. We look forward to the inclusion of additional data into this combined fit from the upcoming

short-baseline near detector and ICARUS experiments in order to shed more light onto the $3 + 1$ sterile neutrino hypothesis [27].

We acknowledge the support of Fermilab, the Department of Energy, and the National Science Foundation, and we acknowledge Los Alamos National Laboratory for LDRD funding. This document was prepared by the MiniBooNE 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.

*Present address: The Aerospace Corporation, Los Angeles, California 90009, USA.

†Deceased.

- [1] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), A Search for Electron Neutrino Appearance at the $\Delta m^2 \sim 1 \text{ eV}^2$ Scale, *Phys. Rev. Lett.* **98**, 231801 (2007).
- [2] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Updated MiniBooNE neutrino oscillation results with increased data and new background studies, *Phys. Rev. D* **103**, 052002 (2021).
- [3] P. A. Zyla *et al.* (Particle Data Group), Review of Particle Physics, *Prog. Theor. Exp. Phys.* **2020**, 083C01 (2020).
- [4] S. Schael *et al.* (ALEPH, DELPHI, L3, OPAL, SLD, LEP Electroweak Working Group, SLD Electroweak Group, SLD Heavy Flavour Group Collaborations), Precision electroweak measurements on the Z resonance, *Phys. Rep.* **427**, 257 (2006).
- [5] V. Brdar and J. Kopp, An Altarelli cocktail for the MiniBooNE anomaly?, *Phys. Rev. D* **105**, 115024 (2022).
- [6] M. Dentler, A. Hernández-Cabezudo, J. Kopp, P. A. N. Machado, M. Maltoni, I. Martínez-Soler, and T. Schwetz, Updated global analysis of neutrino oscillations in the presence of eV-scale sterile neutrinos, *J. High Energy Phys.* **08** (2018) 010.
- [7] J. Kopp, P. A. N. Machado, M. Maltoni, and T. Schwetz, Sterile neutrino oscillations: The global picture, *J. High Energy Phys.* **05** (2013) 050.
- [8] P. B. Denton, Sterile Neutrino Searches with MicroBooNE: Electron Neutrino Disappearance, *Phys. Rev. Lett.* **129**, 061801 (2022).
- [9] C. A. Argüelles, I. Esteban, M. Hostert, K. J. Kelly, J. Kopp, P. A. N. Machado, I. Martínez-Soler, and Y. F. Perez-Gonzalez, MicroBooNE and the ν_e Interpretation of the MiniBooNE Low-Energy Excess, *Phys. Rev. Lett.* **128**, 241802 (2022).
- [10] R. Acciarri *et al.* (MicroBooNE Collaboration), Design and construction of the MicroBooNE detector, *J. Instrum.* **12**, P02017 (2017).
- [11] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for an anomalous excess of charged-current quasi-elastic ν_e interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction, *Phys. Rev. D* **105**, 112003 (2022).
- [12] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for an Excess of Electron Neutrino Interactions in MicroBooNE

- Using Multiple Final State Topologies, *Phys. Rev. Lett.* **128**, 241801 (2022).
- [13] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for an anomalous excess of charged-current ν_e interactions without pions in the final state with the MicroBooNE experiment, *Phys. Rev. D* **105**, 112004 (2022).
- [14] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for an anomalous excess of inclusive charged-current ν_e interactions in the MicroBooNE experiment using wire-cell reconstruction, *Phys. Rev. D* **105**, 112005 (2022).
- [15] MicroBooNE Collaboration, Search for an anomalous excess of charged-current quasi-elastic ν_e interactions with the MicroBooNE experiment using deep-learning-based reconstruction, HEPData (collection) (2021), [10.17182/hepdata.114859](https://doi.org/10.17182/hepdata.114859).
- [16] MicroBooNE Collaboration, Search for an anomalous excess of inclusive charged-current ν_e interactions in the MicroBooNE experiment using wire-cell reconstruction HEPData (collection) (2022), [10.17182/hepdata.114862](https://doi.org/10.17182/hepdata.114862).
- [17] C. A. Argüelles, A. Schneider, and T. Yuan, A binned likelihood for stochastic models, *J. High Energy Phys.* **06** (2019) 030.
- [18] Z. Hu, J. Ling, J. Tang, and T. Wang, Global oscillation data analysis on the 3ν mixing without unitarity, *J. High Energy Phys.* **01** (2021) 124.
- [19] S. A. R. Ellis, K. J. Kelly, and S. W. Li, Current and future neutrino oscillation constraints on leptonic unitarity, *J. High Energy Phys.* **12** (2020) 068.
- [20] S. Parke and M. Ross-Lonergan, Unitarity and the three flavor neutrino mixing matrix, *Phys. Rev. D* **93**, 113009 (2016), [arXiv:1508.05095](https://arxiv.org/abs/1508.05095).
- [21] V. V. Barinov *et al.*, Results from the Baksan Experiment on Sterile Transitions (BEST), *Phys. Rev. Lett.* **128**, 232501 (2022).
- [22] K. Goldhagen, M. Maltoni, S. Reichard, and T. Schwetz, Testing sterile neutrino mixing with present and future solar neutrino data, *Eur. Phys. J. C* **82**, 116 (2022).
- [23] J. M. Berryman and P. Huber, Sterile neutrinos and the global reactor antineutrino dataset, *J. High Energy Phys.* **01** (2021) 167.
- [24] F. Kaether, W. Hampel, G. Heusser, J. Kiko, and T. Kirsten, Reanalysis of the GALLEX solar neutrino flux and source experiments, *Phys. Lett. B* **685**, 47 (2010).
- [25] J. N. Abdurashitov *et al.* (SAGE Collaboration), Measurement of the solar neutrino capture rate with gallium metal. III: Results for the 2002–2007 data-taking period, *Phys. Rev. C* **80**, 015807 (2009).
- [26] S. S. Wilks, The large-sample distribution of the likelihood ratio for testing composite hypotheses, *Ann. Math. Stat.* **9**, 60 (1938).
- [27] M. Antonello *et al.* (MicroBooNE, LAr1-ND, ICARUS-WA104 Collaborations), A proposal for a three detector short-baseline neutrino oscillation program in the Fermilab booster neutrino beam, [arXiv:1503.01520](https://arxiv.org/abs/1503.01520).
- [28] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.129.201801> for likelihood construction, MicroBooNE inclusive analysis, and MicroBooNE-only constraints.

Supplemental Materials: MiniBooNE and MicroBooNE Combined Fit to a 3+1 Sterile Neutrino Scenario

S1. LIKELIHOOD

The physics parameters of the model are the mass squared splitting Δm^2 , electron-sterile mixing $|U_{e4}|^2$, and muon-sterile mixing $|U_{\mu 4}|^2$. The mixing parameters ($|U_{e4}|^2$, $|U_{\mu 4}|^2$) are allowed to vary between 0 and 1 while maintaining unitarity of the mixing matrix through the condition $|U_{e4}|^2 + |U_{\mu 4}|^2 \leq 1$. The additional nuisance parameters of the model are the MicroBooNE per-bin systematic scalings α_i . Here the set of physics parameters are denoted by $\vec{\theta}$, and the set of nuisance parameters denoted by $\vec{\eta}$. The combined MiniBooNE-MicroBooNE likelihood is the product the two experimental likelihoods such that

$$\mathcal{L}(\vec{\theta}, \vec{\eta} | \vec{x}) = \mathcal{L}_{\text{MB}}(\vec{\theta} | \vec{x}_{\text{MB}}) \times \mathcal{L}_{\text{uB}}(\vec{\theta}, \vec{\eta} | \vec{x}_{\text{uB}}),$$

where \mathcal{L}_{MB} is the MiniBooNE likelihood, \mathcal{L}_{uB} is the MicroBooNE likelihood, \vec{x}_{MB} is collection of the MiniBooNE data counts, \vec{x}_{uB} is the collection of MicroBooNE data counts, and $\vec{x} = \vec{x}_{\text{MB}} \cup \vec{x}_{\text{uB}}$ is the collection of all data counts. The MiniBooNE likelihood is approximated as a multivariate normal distribution

$$\mathcal{L}_{\text{MB}}(\vec{\theta}, \vec{\eta} | \vec{x}_{\text{MB}}) = \mathcal{N}(\vec{x}_{\text{MB}} | \vec{\mu}_{\text{MB}}(\vec{\theta}), \mathbf{\Sigma}_{\text{MB}}(\vec{\theta})),$$

where $\vec{\mu}_{\text{MB}}$ is the predicted number of data counts in each bin, and $\mathbf{\Sigma}_{\text{MB}}$ is the MiniBooNE covariance matrix. In this case the MiniBooNE covariance matrix includes systematic errors, Poisson statistical errors, and Monte-Carlo statistical errors. The MicroBooNE likelihood is given by

$$\mathcal{L}_{\text{uB}}(\vec{\theta}, \vec{\eta} | \vec{x}_{\text{uB}}) = \mathcal{N}(\vec{\alpha} | 1, \mathbf{\Sigma}_{\text{uB}}) \times \prod_i \mathcal{L}^{\text{Eff}}(\alpha_i \mu_i^{\text{uB}}(\vec{\theta}), \sigma_{i,\text{mc}}^2(\vec{\theta}, \alpha_i) | x_{i,\text{uB}}),$$

where $\mathcal{N}(\vec{\alpha} | 1, \mathbf{\Sigma}_{\text{uB}})$ is the multivariate normal prior on the MicroBooNE systematics scalings, $\mathbf{\Sigma}_{\text{uB}}$ is the MicroBooNE fractional covariance matrix, μ_i is the predicted number of data counts in each bin before systematic modifications, and $\sigma_{i,\text{mc}}^2$ is the Monte-Carlo statistical error on the per-bin data count prediction after the systematics scalings have been applied. The MicroBooNE fractional covariance matrix, $\mathbf{\Sigma}_{\text{uB}}$, is the constrained fractional covariance matrix from [1, 2]. The likelihood \mathcal{L}^{Eff} is a Poisson-based likelihood that accounts for finite Monte-Carlo sample errors, and is described in [3]. The minimum $-\log \mathcal{L}$ for each of the 3+1 fit scenarios described in the main text is given in Table S1.

Datasets Included	$\min_{\vec{\eta}} \{-\log \mathcal{L}\}$ (no-oscillation model)	$\min_{\vec{\theta}, \vec{\eta}} \{-\log \mathcal{L}\}$ (3+1 model)
MiniBooNE only	219.1	205.2
MiniBooNE + DL	251.0	238.7
MiniBooNE + WC	673.1	664.5

TABLE S1. Summary of the minimum negative-log-likelihood values for each of the 3+1 fit scenarios described in the main text, including the result for both the no-oscillation and oscillation cases.

S2. COMBINED FIT WITH MICROBOONE INCLUSIVE ANALYSIS

In this section, we provide the predicted event event rate in the MiniBooNE $\nu_e + \bar{\nu}_e$ and $\nu_\mu + \bar{\nu}_\mu$ distributions, as well as well as the MicroBooNE ν_e FC, ν_e PC, ν_μ FC and ν_μ PC distributions for the ‘‘Combination (WC)’’ best fit to the 3+1 model in Table I. Comparisons between data and prediction in each of these channels is shown in Figure S1. One can see that the best-fit solution prefers negligible ν_μ disappearance while still allowing for enough $\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_e(\bar{\nu}_e)$ appearance to explain most of the excess in the MiniBooNE $\nu_e + \bar{\nu}_e$ channel. Additionally, the systematic pull terms in the MicroBooNE analysis modify the prediction in the ν_μ FC and ν_μ PC channels to match the data, in contrast to the central value prediction shown in Figure 21 of Ref. [4].

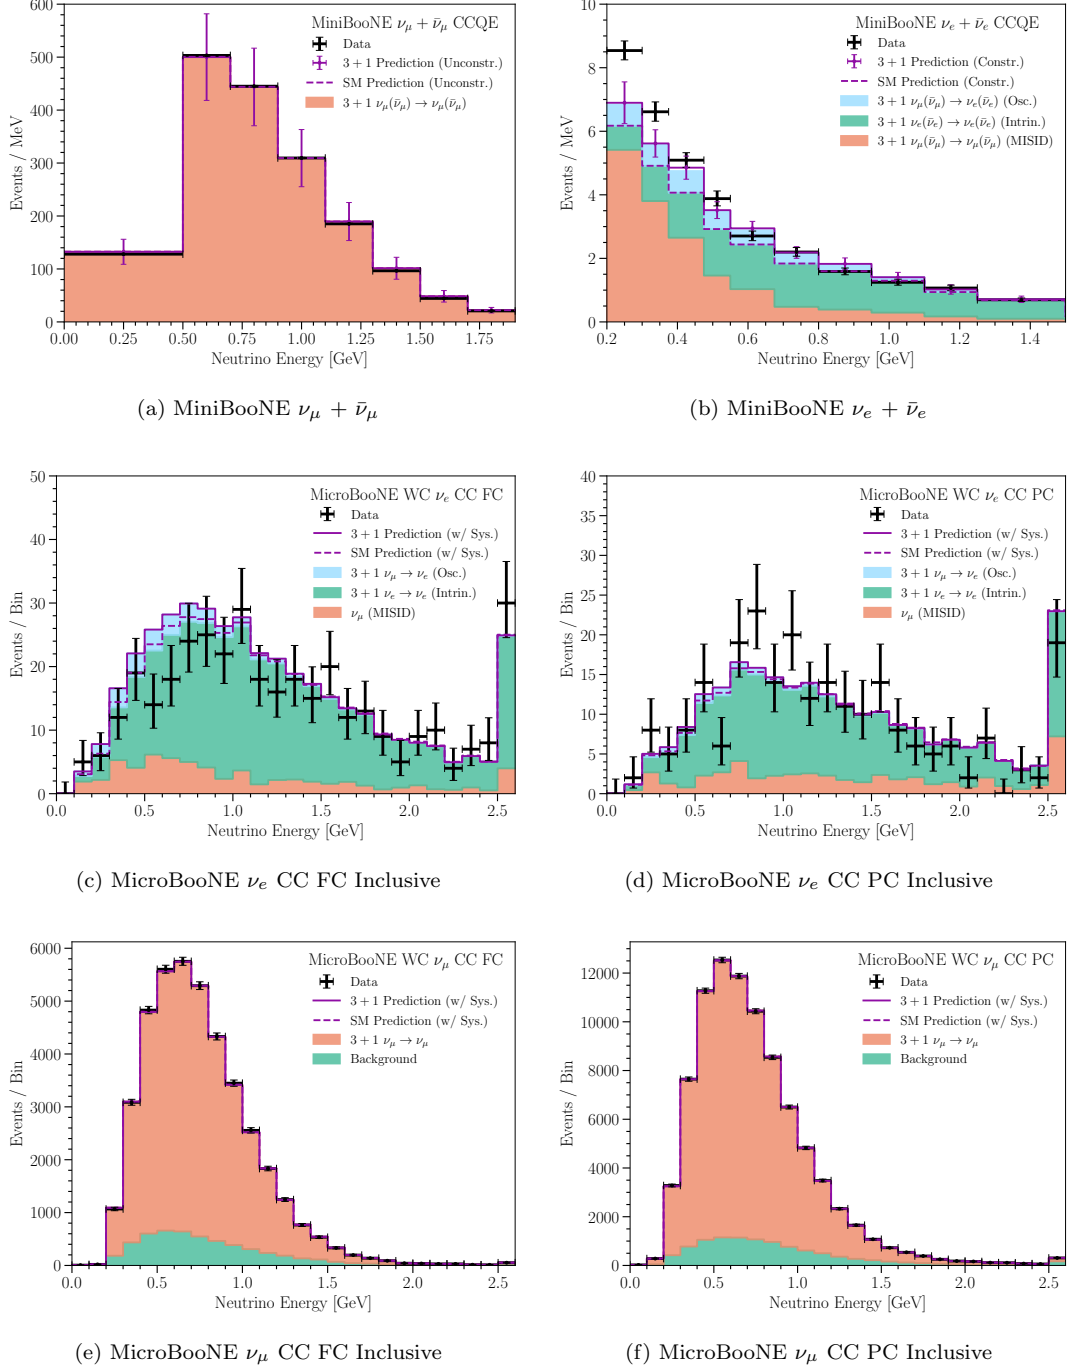


FIG. S1. Comparison between data and prediction for each experiment, showing the prediction from both the SM and the 3+1 model, considering the 3+1 “Combination (WC)” fit parameters of Table I. See the caption of Figure 1 in the main text for a description of the content of each subfigure.

S3. CONSTRAINTS FROM MICROBOONE DATA

In this section, we report the constraints in 3+1 parameter space derived from each of MicroBooNE samples individually. Results from the CCQE (Inclusive) sample are shown in the left (right) plot of Figure S2. One can see that the Inclusive sample sets a slightly stronger constraint than the CCQE sample.

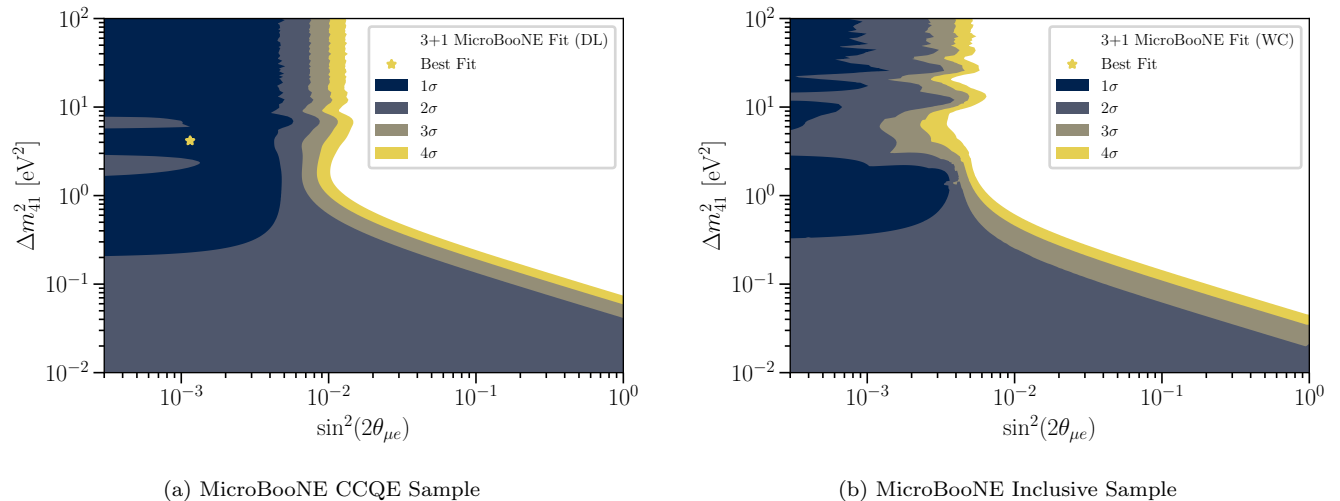


FIG. S2. Constraints in 3+1 parameter space from the MicroBooNE CCQE (left) and Inclusive (right) samples. See the caption of Figure 2 in the main text for more details on the content shown in each plot. Note that the best fit point for the Inclusive sample lies outside of the window shown here.

-
- [S1] P. Abratenko *et al.* (MicroBooNE), Search for an anomalous excess of charged-current quasi-elastic ν_e interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction, (2021), [arXiv:2110.14080 \[hep-ex\]](#).
 - [S2] 'NuE background constrained fractional covariance matrix' of 'Search for an anomalous excess of charged-current quasi-elastic ν_e interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction' (2021).
 - [S3] C. A. Argüelles, A. Schneider, and T. Yuan, A binned likelihood for stochastic models, *JHEP* **06**, 030, [arXiv:1901.04645 \[physics.data-an\]](#).
 - [S4] P. Abratenko *et al.* (MicroBooNE), Search for an anomalous excess of inclusive charged-current ν_e interactions in the MicroBooNE experiment using Wire-Cell reconstruction, (2021), [arXiv:2110.13978 \[hep-ex\]](#).

5.3.2 Publication: *Implications of MicroBooNEs low sensitivity to electron antineutrino interactions in the search for the MiniBooNE excess*

The first MicroBooNE results indeed set strong limits on ν_e CC interactions as an explanation for the entire MiniBooNE LEE. However, as pointed out in Ref. [201], this conclusion changes if the MiniBooNE LEE comes instead from $\bar{\nu}_e$ CC interactions. This is driven mainly by differences in the low-energy suppression of $\bar{\nu}_e$ -nucleus and ν_e -nucleus cross sections, coming from interference between the vector and axial form factors as well as nucleon binding energy effects. Binding energy effects are especially important in MicroBooNE—within a non-isoscalar nucleus like ^{40}Ar , it takes more energy for a $\bar{\nu}_e$ to turn a proton into a neutron than it does for a ν_e to turn a neutron into a proton. This is not necessarily the case in MiniBooNE, as ^{12}C is an isoscalar nucleus. This study was led by myself and Matheus Hostert. The full *Physical Review D* publication can be found below. The most important results are shown in Figure 5 and Figure 6, which indicate that the MicroBooNE sensitivity to the MiniBooNE LEE decreases as more of the excess is attributed to $\bar{\nu}_e$ rather than ν_e CC interactions. Specifically, the MicroBooNE data are consistent at the 2σ confidence level with a scenario in which the entire MiniBooNE LEE comes from $\bar{\nu}_e$ interactions.

Implications of MicroBooNE's low sensitivity to electron antineutrino interactions in the search for the MiniBooNE excess

N. W. Kamp¹, M. Hostert^{2,3,4}, C. A. Argüelles⁵, J. M. Conrad¹, and M. H. Shaevitz⁶

¹Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

²Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2J 2W9, Canada

³School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA

⁴William I. Fine Theoretical Physics Institute, School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA

⁵Department of Physics and Laboratory for Particle Physics and Cosmology, Harvard University, Cambridge, Massachusetts 02138, USA

⁶Columbia University, New York, New York 10027, USA



(Received 10 February 2023; accepted 12 April 2023; published 3 May 2023)

The MicroBooNE experiment searched for an excess of electron-neutrinos in the Booster Neutrino Beam (BNB), providing direct constraints on ν_e -interpretations of the MiniBooNE low-energy excess (LEE). In this article, we show that if the MiniBooNE LEE is caused instead by an excess of $\bar{\nu}_e$, then liquid argon detectors, such as MicroBooNE, SBND, and ICARUS, would have poor sensitivity to it. This is due to a strong suppression of $\bar{\nu}_e$ - ^{40}Ar cross sections in the low-energy region of the excess. The MicroBooNE results are consistent at the 2σ CL with a scenario in which the MiniBooNE excess is sourced entirely by $\bar{\nu}_e$ interactions. The opportune location of ANNIE, a Gd-loaded water Cherenkov detector, allows for a direct search for a $\bar{\nu}_e$ flux excess in the BNB using inverse beta-decay events.

DOI: [10.1103/PhysRevD.107.092002](https://doi.org/10.1103/PhysRevD.107.092002)

I. INTRODUCTION

The MiniBooNE Experiment at Fermi National Accelerator Laboratory (Fermilab) used a 450 t fiducial volume mineral-oil-based (CH_2) Cherenkov detector to search for the appearance of electronlike events in a beam made predominantly of muon-flavor neutrinos. The beam, produced in the Booster Neutrino Beamline (BNB), resulted from 8.9 GeV total energy protons impinging on a beryllium target, with charged mesons magnetically focused toward the detector [1]. The polarity of the magnet could be switched to allow either positively or negatively charged mesons to be focused. The pions and kaons decayed to mainly produce a muon-flavor flux of neutrinos or antineutrinos, with low electron-flavor content, as discussed below. The beam traversed largely undisturbed to reach the MiniBooNE detector located 541 m downstream. The analysis sought to isolate the charged-current quasielastic (CCQE) scattering of neutrinos, $\nu_e + n \rightarrow e^- + p$, or antineutrinos, $\bar{\nu}_e + p \rightarrow e^+ + n$.

In the MiniBooNE Cherenkov detector, both reactions appear as single electromagneticlike Cherenkov rings.

During a series of runs from 2002 to 2019, the MiniBooNE experiment received 18.75×10^{20} (11.27×10^{20}) protons on target (POT) with the magnetic horn focusing positively (negatively) charged mesons. An excess of 560.6 ± 119.6 (77.4 ± 28.5) electronlike events above the background from intrinsic electron-flavor flux and misidentified muon-flavor events was observed [2]. These low-energy excesses are often referred to as the MiniBooNE “LEE” signal. The community has engaged in a thorough search for misidentified particles that were not included in the analysis but has not identified a conclusive explanation behind the full excess [3–5]. This leads us to reconsider the flux that could cause the LEE signature.

Because the detector is limited to identifying an electronlike Cherenkov ring, it is not possible to identify the neutrino versus antineutrino content of the LEE. Therefore, under the assumption that the LEE is caused by an excess of ν_e or $\bar{\nu}_e$ CCQE events in the detector, we can categorize the possible explanations as follows (defining $f_{\bar{\nu}_e/\nu_e}$ to be the $\bar{\nu}_e$ fractional contribution to the LEE in neutrino mode):

- (1) *Scenario 1*: The excess is entirely due to ν_e interactions in neutrino mode running, $f_{\bar{\nu}_e/\nu_e} = 0$, and entirely $\bar{\nu}_e$ in antineutrino mode running. This fits the classic model of sterile-enhanced $\nu_\mu \rightarrow \nu_e$ neutrino oscillations.

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.

- (2) *Scenario 2*: The excess arises from a flux of mixed content. In this scenario, we assume the event rate of the excess is evenly split into neutrino and anti-neutrino events, $f_{\bar{\nu}/\nu} = 0.5$. In neutrino mode, this hypothesis corresponds to a flux excess of antineutrinos that is larger and lower-energy than the flux excess of neutrinos.
- (3) *Scenario 3*: The excess is entirely due to $\bar{\nu}_e$ interactions in neutrino and antineutrino mode running, $f_{\bar{\nu}/\nu} = 1$.

Of these three possibilities, only *Scenario 1* has been thoroughly explored by the community [5]. While the two experiments were situated in very different beams, we note that an anomalous flux of antineutrinos in MiniBooNE may also be compatible with the unexplained signal at the LSND experiment. LSND operated with a liquid scintillator detector at the LANSCE spallation source and observed a 3.8σ -significant excess of inverse beta decay (IBD) [6]. Because of the unique IBD signature, a positron accompanied by delayed neutron capture, the LSND excess favors a $\bar{\nu}_e$ interpretation over a ν_e one.

Motivated by *Scenario 1*, the MicroBooNE Experiment was proposed to run on the same Booster Neutrino Beamline, 70 m upstream of the MiniBooNE detector, using an 80 t fiducial volume liquid argon time projection chamber (LArTPC). The LArTPC technology was selected in 2006 to greatly reduce photon-electron misidentification backgrounds that were thought, at the time, to be the best explanation of the MiniBooNE anomaly [7]. However, as a state-of-the-art detector, the cost per ton for the detector led to a restricted size, and hence low statistics. The experiment has published data taken from 2016–2018 totaling 6.9×10^{20} POT in neutrino mode. In principle, the $\mathcal{O}(1)$ cm vertex resolution of the MicroBooNE detector [8] and its ability to detect protons allows for the separation of $\nu_e + n \rightarrow e^- + p$ and $\bar{\nu}_e + p \rightarrow e^+ + n$ events, making it ideal for testing mixed models like *Scenarios 2* and *3*. Unfortunately, in practice, the argon target has a highly-suppressed antineutrino interaction cross section in the $E_{\bar{\nu}} < 600$ MeV range of interest for the LEE, as we explain below. Thus, given the size of the detector and the length of the run, MicroBooNE is much less sensitive to an antineutrino component in the LEE.

MicroBooNE performed a model-agnostic search for the MiniBooNE LEE [9], building a template of the excess with respect to the Standard Model prediction from the MiniBooNE neutrino-mode data. The LEE template was derived by unfolding the difference between the central value of the data and background predictions at MiniBooNE to an excess of neutrinos in the beam. This process assumed that the MiniBooNE LEE signal was entirely due to neutrino interactions. The strategy to obtain the template and corresponding constraint on it at MicroBooNE was the following [10]:

- (1) The MiniBooNE Monte Carlo sample of true ν_e charged-current (CC) interactions is used to construct

the response matrix $\mathcal{A}_{i\alpha} \equiv P(\text{Reconstructed } i | \text{generated in } \alpha)$, where i refers to the reconstructed energy bin $(E_{\nu}^{\text{QE}})_i$, and α refers to the true-energy bin $(E_{\nu}^{\text{true}})_{\alpha}$.

- (2) The unfolded intrinsic ν_e CC interaction rate in true energy space, u_{α} , is obtained via the D’Agostini iterative approach [11]. The regularization parameter was chosen to (1) minimize the variance of the unfolded spectra, (2) minimize the bias of the unfolded spectra, and (3) produce an expected event rate in reconstructed energy space, which is statistically consistent with the observed MiniBooNE data.
- (3) The ratio of the unfolded event rate u_{α} and the central value MiniBooNE Monte Carlo prediction are taken as weights in true neutrino energy space, which are then applied to true ν_e CC interactions in the MicroBooNE simulation to produce the LEE model prediction.

The two MicroBooNE analyses with the highest sensitivity to the LEE were the Deep-Learning Based Analysis (DL) [12] and the Wire-Cell Analysis (WC) [13]. In brief, the DL analysis looked for an exclusive sample of ν_e CC interactions with one electron and one proton in the final state ($1e1p$). The reconstruction chain relied on two novel LArTPC-specific deep learning algorithms, SparseSSNet [14] and MPID [15], to isolate these $1e1p$ events. The kinematics of the electron and proton were required to be consistent with CC quasielastic scattering to reduce systematic uncertainties on the interaction cross-section. This resulted in a signal sample with a large signal-to-background ratio but comparatively low statistics, with 25 events passing the full selection. In the $E_{\nu/\bar{\nu}} < 400$ MeV range, because the DL analysis required a lepton-proton vertex, the selected events were almost entirely due to neutrino interactions rather than antineutrino interactions. The energy distribution of DL-selected events was fit to the Standard Model prediction plus the MiniBooNE-based LEE model with floating normalization. Based on this fit, the DL analysis limited the content of the MiniBooNE LEE to $< 38\%$ ν_e interactions at 2σ . Since the DL analysis is not sensitive to antineutrinos, this result implies that more than 68% of the excess can be due to $\bar{\nu}_e$ or other unrelated non-neutrino events.

The WC analysis looked for an inclusive sample of ν_e CC interactions with one electron and anything else in the final state ($1eX$). The namesake Wire-Cell algorithm [16] was used to identify three-dimensional space points of charge within each MicroBooNE image, which were then clustered and analyzed using a series of pattern recognition algorithms to isolate $1eX$ events. The WC analysis allowed for single-lepton events as well as multiprong vertices. Thus, in the $E_{\nu/\bar{\nu}} < 400$ MeV range, soft- as well as hard-scattering processes dominate. For WC, antineutrino events could contribute, although the missing energy due to the neutron would lead to an underestimate of the antineutrino

energy. Fitting the Standard Model prediction plus the MiniBooNE-based LEE model with floating normalization to the WC selected events, MicroBooNE set a 2σ upper limit of 50.2% electron-flavor content in the LEE [13]. In quoting this result, the Collaboration assumed that the Wire-Cell analysis is not sensitive to the portion of the excess that is caused by sources other than an excess of electron neutrinos.

In the remainder of this article, we explore how the constraints above change when the assumption on the neutrino-antineutrino composition of the LEE is allowed to vary. While the design of the DL analysis makes it insensitive to the content of antineutrinos in the LEE, it can still, in principle, be observed by the WC analysis. As we will show, however, due to the fact that antineutrino-argon cross sections are much smaller than those of antineutrino- CH_2 for the LEE energies, the entire suite of MicroBooNE analyses turns out to be significantly less effective in constraining an excess of antineutrinos. This is shown in Fig. 1, which indicates that the number of excess events predicted in the lowest energy region of the WC analysis decrease as the $\bar{\nu}_e$ content of the LEE increases from *Scenario 1* to *Scenario 3*.

The rest of this article proceeds as follows. First, in Sec. II, we compare the neutrino versus antineutrino interaction rates in the MiniBooNE and MicroBooNE targets, CH_2 and Ar, respectively. We then generalize the unfolding of the MiniBooNE LEE to scenarios with a mixture of neutrino and antineutrino fluxes in Sec. III. We apply our procedure to three assumptions on the

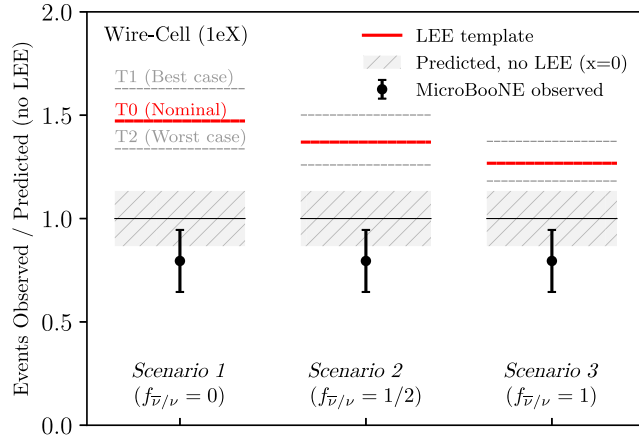


FIG. 1. The ratio between observed and predicted LEE events at MicroBooNE in the three antineutrino-neutrino LEE scenarios. The gray bands show the uncertainty in the prediction in the absence of a LEE ($x = 0$), and in red, we show the prediction for the LEE with ($x = 1$) in the three scenarios. From left to right, we show *Scenario 1* ($f_{\bar{\nu}/\nu} = 0$), *Scenario 2* ($f_{\bar{\nu}/\nu} = \frac{1}{2}$), and *Scenario 3* ($f_{\bar{\nu}/\nu} = 1$). The fainter dashed gray lines represent the LEE prediction for the two alternative templates of Ref. [17]. For the assumptions and methodology behind our analysis, see Sec. III and the Appendix.

neutrino-antineutrino composition of the LEE event rate: 100% ν_e , a 50%–50% mix, and 100% $\bar{\nu}_e$. Finally, we discuss the implications of these scenarios for the MicroBooNE analyses in Sec. IV and discuss the next steps in identifying sources of $\bar{\nu}_e$ in the MiniBooNE LEE in Sec. V.

II. NEUTRINO AND ANTINEUTRINO INTERACTIONS IN THE LEE

In this section, we discuss the interaction cross sections for electron-neutrinos and electron-antineutrinos in the target material of the MiniBooNE and MicroBooNE detectors. Their impact in the MicroBooNE analyses is discussed in Sec. IV.

At the fundamental level, the interaction cross sections of neutrinos and antineutrinos are different due to the sign of the axial-vector component. For CCQE scattering on free nucleons, the difference can be expressed quite simply as

$$\frac{1}{\sigma_0} \left(\frac{d\sigma^\nu}{dy} - \frac{d\sigma^{\bar{\nu}}}{dy} \right) = y \left(1 - \frac{y}{2} \right) (F_1 + F_2) F_A, \quad (1)$$

where $\sigma_0 \equiv \frac{G_F^2 |V_{ud}|^2 M E_\nu}{\pi}$ and $y = 1 - E_\ell/E_\nu$ is the inelasticity parameter. In addition to the dependence on y , the above expression has an implicit dependence on the kinematics through the nucleon form factors $F_i \equiv F_i(Q^2)$, where $Q^2 = 2E_\nu M y$ is the momentum exchange with the nucleon. Equation (1) constitutes the interference between axial-vector, $F_A \propto g_A$, and vector pieces of the amplitude. This interference is destructive for antineutrinos but constructive for neutrinos. Interestingly, Eq. (1) leads to a preference for lower momentum transfer and thus more forward scattering angles of the final state lepton in the antineutrino case compared to the neutrino case, a behavior that is in better agreement with the forward-peaked nature of the MiniBooNE excess [2].

For isoscalar targets, and in the absence of thresholds, $\sigma_\nu > \sigma_{\bar{\nu}}$. At high energies, the ratio asymptotes to a factor of $\sim 1/2$. At low energies, however, it can vary significantly due to threshold, nuclear, and binding-energy effects.

A. Cross sections at MiniBooNE and MicroBooNE

The composition of the mineral oil in MiniBooNE is CH_2 , providing six bound neutrons for neutrino CC interactions, and six bound and two free protons for antineutrino CC interactions. In MicroBooNE, the argon nuclear targets provide 22 bound neutrons for neutrino CC interactions and 18 bound protons for antineutrinos. These nonisoscalar materials enhance the number of nucleon targets for antineutrino CCQE at MiniBooNE by 33% and suppress them at MicroBooNE by 10%.

However, an even stronger effect is at play in the energy region of the LEE; the separation energy of protons and neutrons inside the argon nucleus. As opposed to the

isoscalar ^{12}C nucleus, the nonisoscalar ^{40}Ar nucleus contains protons that are more strongly bound than neutrons, and, therefore, require more energy to be knocked out by the CC interactions of antineutrinos. This effect is important at small neutrino energies, where the center-of-mass energy is comparable with the nuclear binding energies. In addition, the Pauli blocking of neutrons in argon is more significant than in carbon due to its size.

We show a comparison of the total and exclusive cross sections for neutrino and antineutrino cross sections on carbon and argon in Fig. 2, obtained from GENIE v3.02.00 [18,19]. While this ratio is similar for argon and CH_2 at high energies, it is significantly different at lower energies, varying by factors larger than two in the energy region of the LEE. For the same event rate at MiniBooNE, this implies that MicroBooNE would see fewer events if antineutrinos induced those events rather than neutrinos.

We have elected to use GENIE v3.02.00, as it is the most up-to-date GENIE public release at the time of this study. Other versions do exist—for example, the $\text{CC0}\pi$ MicroBooNE tune of GENIE v3.00.06 presented in Ref. [20]. While it would certainly be interesting to investigate the antineutrino hypothesis within the context of the MicroBooNE GENIE tune, it is not publicly available to our knowledge. Therefore, we rely on GENIE v3.02.00 for this study and leave the consideration of alternative neutrino event generators to future work.

To quantify the effect of an antineutrino component in the LEE, we should also consider the kinematics of the neutrino- and antineutrino-induced leptons. Since the interference term between vector and axial components, shown in Eq. (1) for CCQE, is proportional to the inelasticity parameter $y = 1 - E_e/E_\nu$, when it contributes constructively, it leads to a preference for larger y , and, therefore, lower-energy leptons. This is the case for neutrino-induced reactions. In the case of antineutrinos, the interference is destructive, leading to a preference for smaller y , and, therefore, higher-energy leptons. As we show below, in the context of the LEE, this implies that to reproduce the observed excess of events in *Scenarios 2* and *3*, the flux excess of antineutrinos would require a lower mean energy than the corresponding flux excess of neutrinos in *Scenario 1*.

Note that this also implies that the impact of nuclear physics on the total cross sections, especially the suppression of low- Q^2 configurations, is different between neutrinos and antineutrinos. This dependence on the lepton kinematics and hadronic energy means that the unfolding procedure adopted by MicroBooNE ought to be modified before applying it to the antineutrino hypothesis.

III. THE MiniBooNE UNFOLDING-BASED TEMPLATE ANALYSIS

To unfold the MiniBooNE excess under antineutrino-based explanations of the LEE, we follow the procedure

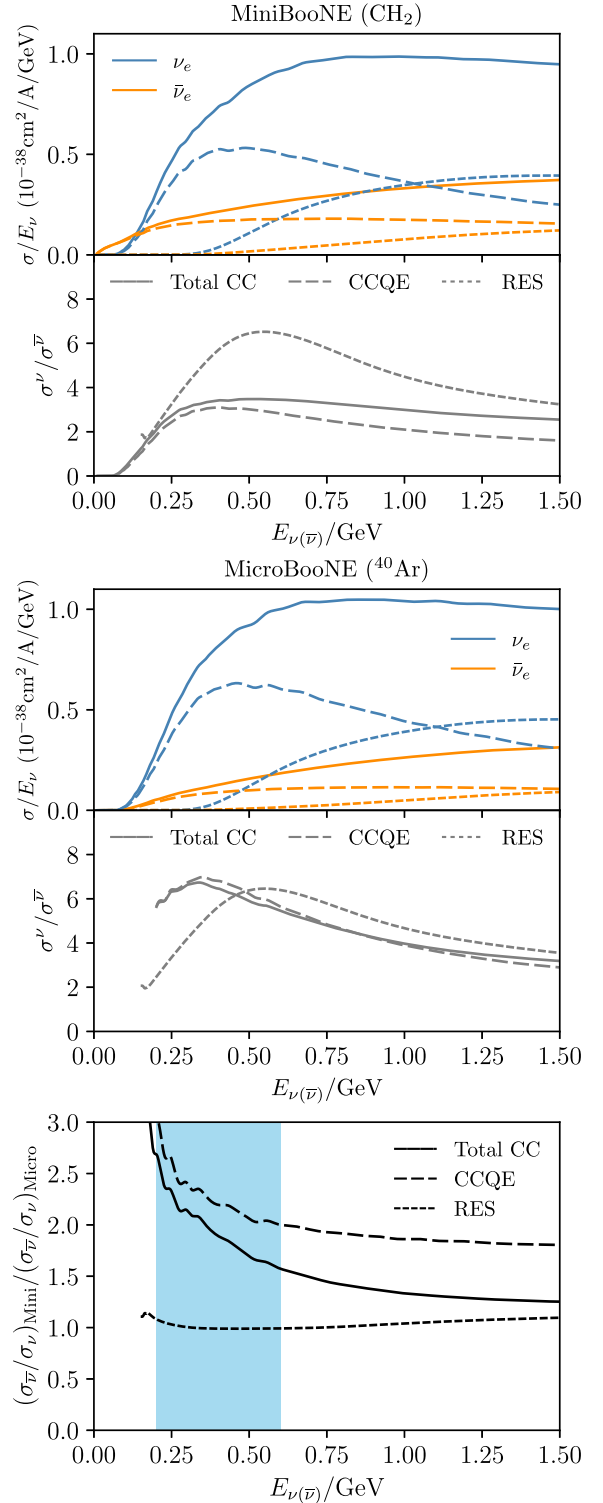


FIG. 2. The top and middle panels show the total ν_e and $\bar{\nu}_e$ cross sections and their ratios for CH_2 and ^{40}Ar , respectively. The CCQE and resonant cross sections are shown as dashed lines. In the bottom panel, we show the ratio between the MiniBooNE and MicroBooNE $\nu/\bar{\nu}$ ratios. The blue region indicates the region of the MiniBooNE LEE.

outlined by the MicroBooNE Collaboration in Ref. [10]. Specifically, we use the D'Agostini iterative approach [11] to unfold the observed MiniBooNE data, where the prediction u_α^k in true energy bin α at iteration k is given by

$$u_\alpha^k = \sum_{i=1}^{n_r} M_{i\alpha}^{k-1} d_i, \quad (2)$$

where the sum goes over each of the n_r reconstructed energy bins, and d_i denotes the observed data in reconstructed bin i . The matrix $M_{i\alpha}^k$ is defined by

$$M_{i\alpha}^k = \frac{A_{i\alpha} u_\alpha^k}{\epsilon_\alpha \sum_{\beta=1}^{n_i} A_{i\beta} u_\beta^k}, \quad (3)$$

where the response matrix $A_{i\alpha}$ is given by

$$A_{i\alpha} = P(\text{reconstructed in } i | \text{generated in } \alpha), \quad (4)$$

and $\epsilon_\alpha \equiv \sum_i A_{i\alpha}$ is the reconstruction efficiency in the true-energy bin α .

A. Introducing an antineutrino component

The neutrino and antineutrino response matrices will differ in MiniBooNE, as the $d\sigma/dy$ distribution prefers smaller (larger) values of y for neutrino (antineutrino) CCQE scattering. We calculate the antineutrino response matrix in MiniBooNE by simulating $\bar{\nu}_e$ charged-current interactions in CH_2 using GENIE v3.02.00. This formalism is unable to account for the MiniBooNE reconstruction efficiency; thus, we instead estimate $R_{i\alpha} \equiv A_{i\alpha}/\epsilon_\alpha$.

As Cherenkov detectors are only sensitive to the final state lepton, MiniBooNE uses E_ν^{QE} ($E_{\bar{\nu}}^{\text{QE}}$) to reconstruct ν_e ($\bar{\nu}_e$) energies, given by the expressions

$$E_\nu^{\text{QE}} = \frac{2(M'_n)E_\ell - ((M'_n)^2 + m_\ell^2 - M_p^2)}{2[(M'_n) - E_\ell + \sqrt{E_\ell^2 - m_\ell^2 \cos^2 \theta_\ell}]}, \quad (5)$$

$$E_{\bar{\nu}}^{\text{QE}} = \frac{2(M'_p)E_\ell - ((M'_p)^2 + m_\ell^2 - M_n^2)}{2[(M'_p) - E_\ell + \sqrt{E_\ell^2 - m_\ell^2 \cos^2 \theta_\ell}]}. \quad (6)$$

Here, M_n and M_p are the neutron and proton mass, E_ℓ , m_ℓ , and θ_ℓ are the lepton energy, mass, and scattering angle, and $M'_{(n/p)} \equiv M_{(n/p)} - E_B$, where the nucleon binding energy E_B is fixed to 34 (30) MeV for neutrinos (antineutrinos). It is important to emphasize that no matter the origin of the underlying event (ν_e or $\bar{\nu}_e$), interactions are always reconstructed using E_ν^{QE} ($E_{\bar{\nu}}^{\text{QE}}$) for data taken in neutrino (antineutrino) mode.

We consider only neutrino mode data in this study, as this is directly comparable to the MicroBooNE neutrino mode data. We approximate $R_{i\alpha}$ by marginalizing over $E_\nu^{\text{QE, true}}$,

using GENIE v3.02.00 to generate the truth-level final-state kinematic distributions of the e^+ which appear in Eq. (5). The details of this calculation are given in Appendix A 1. We separately approximate the $\bar{\nu}_e$ detection efficiency in MiniBooNE by using the provided detection efficiency as a function of electron energy [21]. The details of the efficiency calculation are given in Appendix A 2.

Armed with our calculation for $R_{i\alpha} = A_{i\alpha}/\epsilon_\alpha$ and ϵ_α , we can perform the unfolding procedure using Eq. (3). This produces a prediction for the antineutrino interaction rate in MiniBooNE as a function of true antineutrino energy, hereafter denoted u_α^{MB} , which we will use to predict a signal in MicroBooNE. In the top panel of Fig. 3, we show our unfolded $\bar{\nu}_e$ prediction considering the first 6.46×10^{20} POT of MiniBooNE data—the same dataset used in Ref. [10]. One can see that the unfolded $\bar{\nu}_e$ template peaks at lower (anti)neutrino energy compared with the unfolded ν_e template. In the bottom panel of Fig. 3, we refold the unfolded excess templates back through the MiniBooNE reconstruction. Both the neutrino and antineutrino

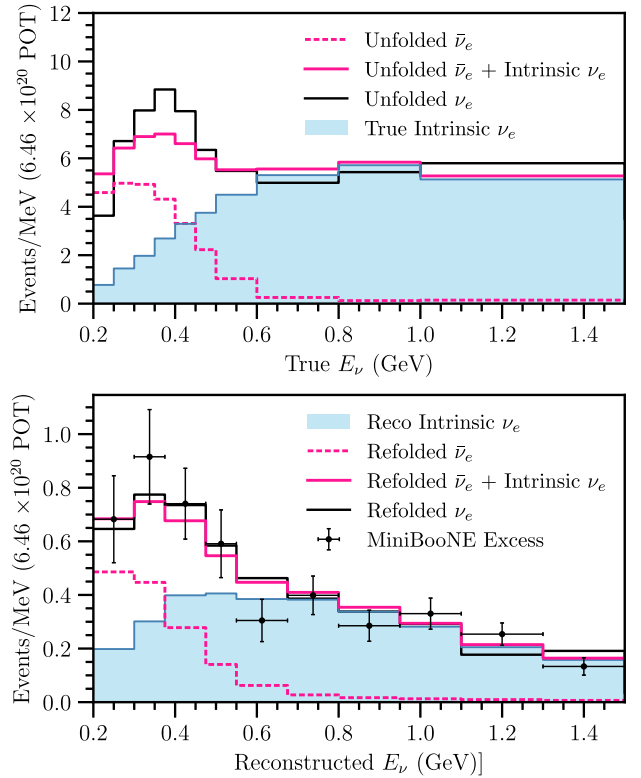


FIG. 3. In the top panel we show the intrinsic ν_e background and the unfolded ν_e and $\bar{\nu}_e$ templates obtained via D'Agostini's unfolding method. The ν_e unfolding procedure begins with the intrinsic ν_e event rate as an initial guess, while the $\bar{\nu}_e$ unfolding procedure begins with a flat distribution in the true neutrino energy. In the bottom panel, we show the refolded prediction in MiniBooNE for both the unfolded ν_e hypothesis and unfolded $\bar{\nu}_e$ plus intrinsic ν_e hypothesis, compared with the excess data for 6.46×10^{20} POT [10].

re-folded predictions are in agreement with the MiniBooNE data.

Once we have an unfolded $\bar{\nu}_e$ MiniBooNE prediction u_a^{MB} , we can fold that prediction into MicroBooNE. In contrast with Fig. 3, here we use the 12.84×10^{20} POT neutrino mode dataset presented in MiniBooNE’s 2018 result [22] for the unfolding process. This is the same dataset used by the MicroBooNE Collaboration to calculate their LEE template [13]. Note that the unfolded spectrum represents the *interaction rate* inside MiniBooNE—thus, we must scale by the ratio of cross sections in Ar and CH₂ when going to MicroBooNE. The predicted MicroBooNE event rate μ_i^{MB} in $E_{\bar{\nu}}^{\text{reco}}$ bin i is given by

$$\mu_i^{\text{MB}} = \sum_{\alpha} \epsilon_{\alpha} R_{i\alpha} u_a^{\text{MB}} \frac{\sigma_{\text{Ar}}((E_{\bar{\nu}}^{\text{true}})_{\alpha})}{\sigma_{\text{CH}_2}((E_{\bar{\nu}}^{\text{true}})_{\alpha})}, \quad (7)$$

where ϵ_{α} and $R_{i\alpha}$ now denote the MicroBooNE $\bar{\nu}_e$ detection efficiency and MicroBooNE response matrix, respectively.

The $\bar{\nu}_e$ response matrix calculation in MicroBooNE is similar to the MiniBooNE calculation. The main difference is that MicroBooNE, being a LArTPC, is able to perform a calorimetric energy reconstruction. The reconstructed energy in MicroBooNE is given by $E_{\bar{\nu}}^{\text{Cal}} = \sum_j (T_j^{\text{reco}} + m_j + B_j)$, where T_j^{reco} , m_j , and B_j denote the observed kinetic energy, rest mass, and binding energy associated with the j th reconstructed final-state particle. The binding energy B_i is taken to be 8.6 MeV for protons and zero for everything else. In the case of $\bar{\nu}_e$ CCQE scattering, MicroBooNE will only reconstruct the final state e^+ . The reconstructed energy $E_{\bar{\nu}}^{\text{reco}}$ is then simply $E_{\bar{\nu}}^{\text{Cal, reco}} = T_{e^+}^{\text{reco}} + m_{e^+}$. Note that this will lead to an underestimation bias in the reconstructed $\bar{\nu}_e$ energy due to the invisible neutron. In order to calculate $R_{i\alpha}$ in MicroBooNE, we marginalize over the $T_{e^+}^{\text{reco}}$ distribution generated using GENIE v3.02.00. The details of this calculation are given in Appendix A 1.

Evaluating the detection efficiency of $\bar{\nu}_e$ interactions in MicroBooNE as a function of $E_{\bar{\nu}_e}$ is more complicated than

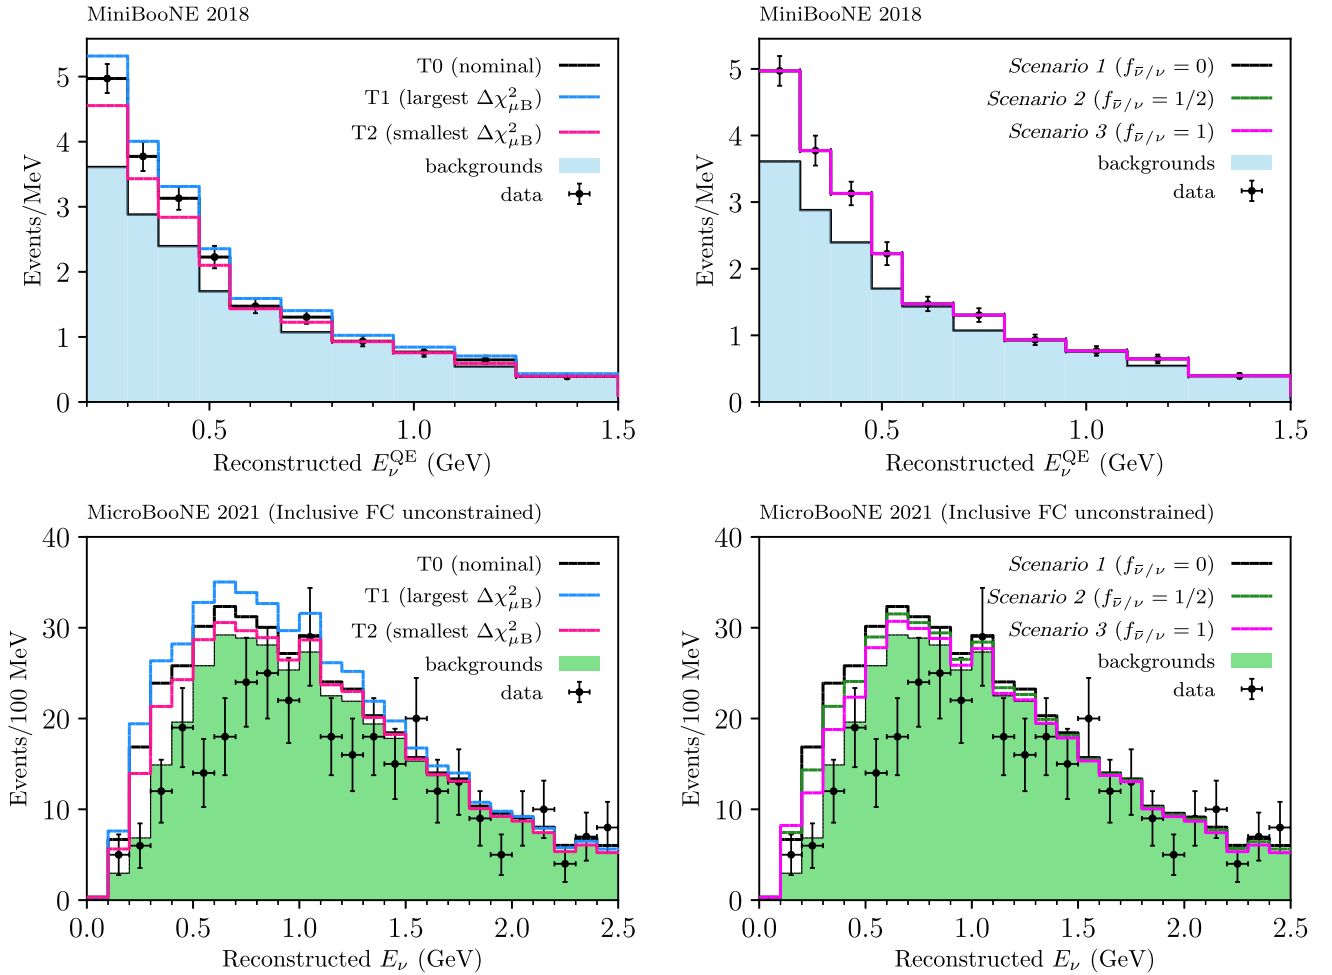


FIG. 4. The MiniBooNE and MicroBooNE spectra in reconstructed neutrino energy. On the left panels, we show three different template choices from Ref. [17] and their corresponding prediction at MicroBooNE, assuming no antineutrinos, $f_{\bar{\nu}/\nu} = 0$. On the right panels, we show how the three different assumptions for the antineutrino composition of the excess impact the nominal template (T0) prediction in MicroBooNE. In MiniBooNE, the three different scenarios of antineutrino compositions look exactly the same, by definition.

the MiniBooNE case. MicroBooNE is not a spherically symmetric detector, thus both the electron direction and energy will impact the detection efficiency. The MicroBooNE reconstruction also relies nontrivially on hadronic information in addition to leptonic information for its nominal ν_e analyses, including the inclusive analysis studied here. Ignoring final state interactions, there will be no hadronic information in $\bar{\nu}_e$ CCQE interactions. It is also possible for ν_e CCQE interactions to occur without hadronic activity if the energy of the final state proton is below the reconstruction threshold; however, this is considerably different than the $\bar{\nu}_e$ case, as the final state neutron can carry away an arbitrary amount of energy without being reconstructed. Given these complications, a detailed estimation of the MicroBooNE $\bar{\nu}_e$ efficiency is out of the scope of this paper. We instead conservatively assume the $\bar{\nu}_e$ efficiency in MicroBooNE to be the same as the reported ν_e efficiency for a given true (anti)neutrino energy. In a realistic scenario, the $\bar{\nu}_e$ efficiency is likely smaller due to the lack of hadronic information; thus, one can interpret the MicroBooNE $\bar{\nu}_e$ prediction derived here as an upper bound.

IV. THE IMPACT ON THE MiniBooNE TEMPLATE ANALYSIS

The right panels of Fig. 4 show the prediction in MicroBooNE under the three different hypotheses for the antineutrino content in the LEE outlined in Sec. I: *Scenario 1* ($f_{\bar{\nu}/\nu} = 0$), *Scenario 2* ($f_{\bar{\nu}/\nu} = 1/2$), and *Scenario 3* ($f_{\bar{\nu}/\nu} = 1$). As can be seen, the LEE template prediction in MicroBooNE from the unfolded MiniBooNE excess decreases in general with $f_{\bar{\nu}/\nu}$, the antineutrino fractional contribution to the excess. Thus, as expected, MicroBooNE is less sensitive to antineutrino-based explanations of the MiniBooNE excess. This is quantified in Fig. 5, which shows the test statistic of the Wire-Cell analysis, $\Delta\chi_{\mu B}^2(x) \equiv \chi_{\mu B}^2(x) - \chi_{\mu B}^2(x=0)$, as a function of the signal strength scaling parameter x introduced in Ref. [9]. As shown in Table I, the exclusion power drops significantly as the predicted antineutrino content becomes larger. Specifically, as $f_{\bar{\nu}/\nu}$ increases from 0 to 0.5, the test statistic $\Delta\chi_{\mu B}^2(x=1.0)$ falls from 13.54 to 8.06. For $f_{\bar{\nu}/\nu} = 1.0$, $\Delta\chi_{\mu B}^2(x=1.0) = 3.82$. At around $x = 0.2$, the $f_{\bar{\nu}/\nu} = 1.0$ case predicts a slightly negative $\Delta\chi_{\mu B}^2$, implying a minor improvement with respect to the nominal BNB prediction. This is most likely caused by the small excess observed in the lowest energy bin of the MicroBooNE analysis, as shown in the lower-right panel of Fig. 4.

Assuming Wilks' theorem [23] with one degree of freedom, the critical $\Delta\chi_{\mu B}^2$ value at the 95.45% (2σ) confidence level is $\Delta\chi_{\mu B}^2 = 4$. We use this to calculate 2σ upper limits on the signal scaling parameter x in *Scenarios 1, 2, and 3*, considering the test statistic




FIG. 5. The MicroBooNE $\Delta\chi^2$ for the inclusive $1eX$ analysis as a function of the signal strength x . The black solid line shows the $\Delta\chi^2$ for the nominal template, following the calculation of Ref. [17]. official MicroBooNE curve, shown in dashed black. In solid green and violet lines, we show the corresponding curves for *Scenarios 2 and 3* for the neutrino-antineutrino fractions of the LEE, respectively.

$\Delta\chi_{\mu B}^2(x) = \chi^2(x) - \min_x \{\chi^2(x)\}$. These 2σ upper limits are shown in Table I. From Table I, it is evident that MicroBooNE is much less sensitive to *Scenario 2* than *Scenario 1*, and is essentially insensitive to *Scenario 3*.

In the left panels of Fig. 4, we show three different templates for the MiniBooNE LEE: T0, T1, and T2. Template T0 is the nominal unfolded template. Templates T1 and T2, defined in Fig. 2 of Ref. [17], correspond to the most and fewest number of excess events obtained in the unfolding procedure while remaining consistent with the MiniBooNE excess at $p > 80\%$. In Fig. 6, we show the $\Delta\chi_{\mu B}^2$ exclusion power of the MicroBooNE data as a function of the antineutrino fraction of the MiniBooNE excess, considering each template separately. Here, $\Delta\chi_{\mu B}^2$ is defined with respect to the nominal BNB prediction, and we assume the nominal signal strength scaling for each template, i.e., $x = 1.0$. From Fig. 6, one can clearly see that $\Delta\chi_{\mu B}^2$ decreases rapidly as the f increases. This is quantified in Table II, which reports the value of $\Delta\chi_{\mu B}^2$ for each template under *Scenario 1* ($f_{\bar{\nu}/\nu} = 0$), *Scenario 2* ($f_{\bar{\nu}/\nu} = 0.5$), and *Scenario 3* ($f_{\bar{\nu}/\nu} = 1.0$). Note that the results for the nominal

TABLE I. Statistical results from the MicroBooNE Wire-Cell analysis on the signal strength scaling parameter x , considering the three different scenarios outlined in Sec. I. $\Delta\chi_{\mu B}^2(x=1.0)$ is reported with respect to the nominal BNB prediction, i.e., without any additional MiniBooNE-like excess.

	$f_{\bar{\nu}/\nu} = 0.0$	$f_{\bar{\nu}/\nu} = 0.5$	$f_{\bar{\nu}/\nu} = 1.0$
$\Delta\chi_{\mu B}^2(x=1.0)$	13.54	8.06	3.82
2σ upper bound on x	0.49	0.69	1.01




FIG. 6. The MicroBooNE $\Delta\chi^2$ for the inclusive $1eX$ analysis as a function of the antineutrino-neutrino fraction of the excess, $f_{\bar{\nu}/\nu}$. In black, we show the variation for the nominal template, while in blue and pink, we repeat the same exercise for the two templates shown in the left panels of Fig. 4.

template (T0) are the same as those reported in Table I. Even for the most optimistic case (T1), the exclusion power is significantly suppressed for $f_{\bar{\nu}/\nu} = 1.0$.

We can define the test statistic $\Delta\chi^2_{\mu B}(f_{\bar{\nu}/\nu}) \equiv \chi^2_{\mu B}(f_{\bar{\nu}/\nu}) - \min_{f_{\bar{\nu}/\nu} \in [0,1]} \{\chi^2_{\mu B}(f_{\bar{\nu}/\nu})\}$ to calculate the 2σ lower limit on $f_{\bar{\nu}/\nu}$ for each template. These lower limits are also reported in Table II; depending on the template they hover around $f_{\bar{\nu}/\nu} \sim 0.5$. Thus if we take the MiniBooNE excess at face value (i.e., restrict to $x = 1.0$) and consider the nominal template (T0), the MicroBooNE results require at least 55% of the MiniBooNE excess come from $\bar{\nu}_e$ events at the 2σ CL. Note that this statement relies on restricting ourselves to $x = 1.0$ —as shown in Table I, $\chi^2_{\mu B}(x = 1.0) - \chi^2_{\mu B}(x = 0) = 8.06$ for $f_{\bar{\nu}/\nu} = 0.5$, meaning that there is just under 3σ tension with the no-excess hypothesis ($x = 0$) when attributing half of the MiniBooNE excess to $\bar{\nu}_e$ events.

TABLE II. Statistical results from the MicroBooNE Wire-Cell analysis on the antineutrino fractional contribution to the MiniBooNE excess $f_{\bar{\nu}/\nu}$, considering the three different templates shown in the left panels of Fig. 4. For the first three rows, $\Delta\chi^2_{\mu B}(f_{\bar{\nu}/\nu})$ is calculated as in Fig. 6, while the last row considers the $\Delta\chi^2_{\mu B}(f_{\bar{\nu}/\nu})$ definition given in the text to calculate 2σ lower limits on $f_{\bar{\nu}/\nu}$, assuming Wilks' theorem for 1 d.o.f.

Template	T0	T1	T2
Using $\Delta\chi^2_{\mu B}(f_{\bar{\nu}/\nu}) = \chi^2_{\mu B}(f_{\bar{\nu}/\nu}, x = 1.0) - \chi^2_{\mu B}(x = 0)$			
$\Delta\chi^2_{\mu B}(f_{\bar{\nu}/\nu} = 1.0)$	3.82	8.52	2.25
$\Delta\chi^2_{\mu B}(f_{\bar{\nu}/\nu} = 0.5)$	8.06	15.74	4.57
$\Delta\chi^2_{\mu B}(f_{\bar{\nu}/\nu} = 0.0)$	13.54	24.63	7.65
Using $\Delta\chi^2_{\mu B}(f_{\bar{\nu}/\nu}) = \chi^2_{\mu B}(f_{\bar{\nu}/\nu}) - \min_{f_{\bar{\nu}/\nu} \in [0,1]} \{\chi^2_{\mu B}(f_{\bar{\nu}/\nu})\}$			
2σ lower bound on $f_{\bar{\nu}/\nu}$	0.55	0.72	0.24

So far, we have focused on the total $1eX$ sample, using the 7-channel fit of Ref. [13]. However, as a consistency check, the Wire-Cell analysis has also performed an 11-channel fit, separating ν_μ and ν_e CC events into samples with and without final-state protons, $0pX\pi$ and $NpX\pi$. Antineutrinos will contribute almost exclusively to the $0pX\pi$ sample, making it a purer sample of $\bar{\nu}_e$ LEE events. The total number of ν_e CC events is approximately even between the two samples, namely 259 $0pX\pi$ and 298 $NpX\pi$ events. Given that the statistical and systematic uncertainties are larger for the $0pX\pi$ sample, and that it does not observe a deficit of events like that of the $NpX\pi$ sample ($N_{\text{data}}/N_{\text{pred}} = (1.00 \pm 0.08 \text{ stat} \pm 0.21 \text{ sys})$ versus $N_{\text{data}}/N_{\text{pred}} = (0.86 \pm 0.06 \text{ stat} \pm 0.17 \text{ sys})$), we do not expect that an 11-channel fit would qualitatively change our conclusions.

V. DISCUSSION

Our main finding is that explanations of the MiniBooNE LEE involving a large contribution of wrong-sign electron-antineutrinos, from new physics or mismodeling in the experimental simulation, remain viable. The relative suppression of antineutrino cross sections in argon, the target material used by MicroBooNE, with respect to CH_2 , the mineral oil target material used by MiniBooNE, means that MicroBooNE is much less sensitive to a low-energy excess of antineutrinos compared to neutrinos. This motivates new strategies to measure the electron-antineutrino component of the BNB.

The templates in the top panel of Fig. 3 unfolded under the hypothesis of an antineutrino-induced LEE indicate a flux excess that is even lower in energy than its neutrino-induced LEE counterpart. As shown in Fig. 7, such an excess would represent a significant deviation from the BNB model prediction for the flux of intrinsic $\bar{\nu}_e$ in neutrino mode [1]. While no source for such hypothetical enhancement has been identified, this study provides a first glimpse into its energy dependence and relative rate. In what follows, we discuss the implication of these findings for a few different antineutrino hypotheses.

A. The BNB model

The BNB flux model [1] predicts that the wrong-sign electron-neutrinos constitute a total of 0.05% (0.2%) of the total flux in neutrino (antineutrino) mode, arising primarily from charged and neutral kaons as well as secondary muons.¹ An antineutrino explanation to the LEE requires a 10 times larger flux of wrong-sign neutrinos than predicted in the BNB model. Considering only the LEE region, $E_\nu < 600$ MeV, it requires a 25 times larger flux.

¹The chirally-suppressed $\pi^- \rightarrow e^- \bar{\nu}_e$ ($\pi^+ \rightarrow e^+ \nu_e$) decays contribute at a much smaller fraction, namely 6.3×10^{-6} (1.6×10^{-5}) of the total neutrino (antineutrino) mode flux.




FIG. 7. A smooth interpolation of the BNB neutrino fluxes in FHC mode as a function of energy. In light blue, we highlight the nominal prediction for the intrinsic $\bar{\nu}_e$ [1]. In orange, we show the required excess $\bar{\nu}_e$ flux to explain the MiniBooNE low-energy excess in *Scenario 3*, using the unfolded template of Fig. 3 (top). The dashed line is the result using the analysis range, $E_\nu > 200$ MeV. The dotted line in the $E_\nu < 200$ MeV region shows an extrapolation of the unfolded template, assuming it follows the same shape as the intrinsic $\bar{\nu}_e$ component.

Below we address the implications of such an enhancement under different hypotheses.

1. Wrong-sign pions

The early decays of forward-going wrong-sign pions in the BNB, $\pi^- \rightarrow \mu^- \bar{\nu}_\mu$, also appears as a peak at $E_\nu \lesssim 400$ MeV. This process is associated with large uncertainties due to the lack of π^- production data in the forward direction [1]. The question then arises—could a large excess of wrong-sign pions in the low-energy region explain the MiniBooNE LEE? The $\pi^- \rightarrow \bar{\nu}_e$ decays cannot be the only source of the excess, as it would be accompanied by an enormous (lower-energy) flux of $\pi^- \rightarrow \bar{\nu}_\mu$, exceeding the total neutrino flux below 600 MeV by more than two orders of magnitude. Another new source of $\bar{\nu}_e$ from wrong-sign pions are secondary muon decays, namely $\pi^- \rightarrow \mu^- \rightarrow \bar{\nu}_e$. The BNB model predicts that neutrinos from secondary muons correspond to a $\sim 0.2\%$ fraction of the neutrinos from the parent pion. To explain the LEE, this would require a ~ 500 times larger $\pi^- \rightarrow \bar{\nu}_\mu$ flux, which is, again, not realistic.

2. Secondary muons

Another logical possibility is that the primary neutrinos from wrong-sign pions are correctly modeled, but the subsequent decays of secondary muons are not. This effect would have to account for a sizeable increase of the average energy of secondary-muon neutrinos, and, more importantly, would require a fiftyfold enhancement of the fraction between μ^- and π^- neutrinos, bringing

it to 10%. Because of the short decay pipeline, it is difficult to conceive of a scenario where so many forward-going muons could contribute to the neutrino flux. A forward-going muon produced at the target has a small probability $P \sim 50 \text{ m}/(\gamma_\mu c \tau_\mu^0) \sim 7.5\%/\gamma_\mu$ of decaying before hitting the beam absorber ~ 50 m downstream. Here, γ_μ is the Lorentz boost of the muon, which can range from $\gamma_\mu = 2$ –10 in the energy region of interest. While the muons can penetrate the absorber and subsequent dirt, they will not all decay to produce a forward-going $\bar{\nu}_e$, so this possibility is also unrealistic.

3. Associated muon-neutrinos

It is also reasonable to assume that a $\bar{\nu}_e$ excess could be accompanied by a ν_μ or $\bar{\nu}_\mu$ excess. The question of whether this muon component can be observed is not straightforward. In principle, such an excess could lead to a higher rate of muons in MiniBooNE, modifying the measured distribution of $\nu_\mu \text{CC}$ events. Because of the different kinematics of neutrinos and antineutrinos, the energy of a $\bar{\nu}_\mu$ could be misreconstructed as a higher-energy ν_μ . In addition, resonant and coherent charged-pion production would also be modified. By studying $\text{CC}\pi^+$ production, MiniBooNE used this method to constrain the wrong-sign $\nu_\mu \text{CCQE}$ events in antineutrino mode [2]. In practice, however, a very low-energy flux excess like that in Fig. 3 (top) would happen in a region close to the kinematic threshold of muon and pion production. This could exacerbate the excess in $\bar{\nu}_e \text{CC}$ events, where threshold effects are not important. Elastic processes like neutrino- and antineutrino-electron scattering would be impacted, but this component makes up less than 2% of the total number of electronlike events observed by MiniBooNE [2]. Finally, we note that when the MicroBooNE Wire-Cell inclusive $1eX$ sample is divided into $0pX\pi$, and $NpX\pi$ events, a small excess of $\nu_\mu \text{CC}$ $0pX\pi$ events is observed in the energy region of 0.3–1.1 GeV [13]. Because $\nu_\mu \text{CC}$ $NpX\pi$ events are in good agreement with the Monte Carlo, this effect, if it grows in significance, could be explained due to an excess of $\bar{\nu}_\mu$ in the BNB.

To conclude this section, we have found that the full $\bar{\nu}_e$ excess shown Fig. 7 would require a significant deviation from the BNB model presented in Ref. [1]. We remind the reader that our discussion focused on neutrino mode only. Since the origin of the excess in neutrino and antineutrino modes are both unknown, the neutrino-to-antineutrino ratio of the excess in each mode is, in principle, not necessarily the same. For that reason, we do not derive constraints on the excess in antineutrino mode using MicroBooNE data, as the latter was obtained with the beam in neutrino mode. We leave a detailed study of possible anomalous sources of wrong-sign neutrinos in each mode to future work.

B. New physics

We now comment on a few possibilities for a beyond-the-standard-model (BSM) origin of the $\bar{\nu}_e$ excess in MiniBooNE. Neutrino-antineutrino oscillations $\nu_\mu \rightarrow \bar{\nu}_e$ can convert left-handed neutrinos into right-handed antineutrinos [24]. However, they require a chirality flip and are usually too small to be observable due to the smallness of m_ν^2/E_ν^2 . Large neutrino magnetic moments in strong magnetic fields can also induce such oscillations through spin-flip precession [25–27] but are not relevant for short-baseline experiments. In general, $\nu \rightarrow \bar{\nu}$ oscillations are constrained experimentally by solar neutrino experiments [28–30] and by direct searches [31].

A $\bar{\nu}_e$ excess could also stem from exotic pion decays at the target. For instance, the lepton-flavor- and lepton-number-violating branching ratio for the pion, $\pi^+ \rightarrow \mu^+ \bar{\nu}_e$, could produce $\bar{\nu}_e$ with the same energy spectrum as the neutrinos from $\pi^+ \rightarrow \mu^+ \nu_\mu$. The best experimental limits on this decay come from the BEBC detector, which sat in the wide-band high-energy neutrino beam at CERN [31]. A dedicated search finds $\mathcal{B}(\pi^+ \rightarrow \mu^+ \bar{\nu}_e) < 0.15\%$ and $\mathcal{B}(\pi^+ \rightarrow \mu^+ \nu_e) < 0.8\%$ at 90% CL. Precision tests of lepton-flavor universality are also sensitive to this decay channel. Comparing the SM prediction [32] with the experimental measurements [33–36] of $R_{e/\mu} = \Gamma(\pi^+ \rightarrow e^+ \nu)/\Gamma(\pi^+ \rightarrow \mu^+ \nu)$, we find

$$\frac{\Gamma(\pi^+ \rightarrow \mu^+ \bar{\nu}_e)}{\Gamma(\pi^+ \rightarrow \mu^+ \nu_\mu)} = (0.20 \pm 0.19)\%, \quad (8)$$

providing stronger limits on the branching ratio, $\mathcal{B}(\pi^+ \rightarrow \mu^+ \bar{\nu}_e) < 0.50\%$ at 90% CL. We note that this branching ratio can also be constrained by the LSND experiment using low-energy IBD events; in the energy of interest, $E_{e^+} < E_{\bar{\nu}_e}^{\text{DAR}} \sim 30$ MeV, good agreement is found between data and Monte Carlo. We leave a careful evaluation of this limit for future work.

Light particles produced in the beam could be another source of antineutrinos. Decaying light sterile neutrinos can produce antineutrinos in two scenarios. If the sterile neutrino is a Majorana particle, antineutrinos can be produced in the decay $\nu_4 \rightarrow \bar{\nu} \phi$ [37–40], where ϕ is a neutrinophilic light scalar particle. If the sterile neutrino is a Dirac particle, then the subsequent decays of the scalar particle, $\nu_4 \rightarrow \nu(\phi \rightarrow \bar{\nu} \nu)$, can produce antineutrinos at a fraction of 1:2 [38,39]. Finally, lepton-number-charged scalars, ϕ_2 , can lead to antineutrinos in the decay of $\nu_4 \rightarrow \bar{\nu} \phi_2$. In detection, the emission of ϕ_2 can lead to what is effectively an off shell antineutrino scattering process, $\nu_\mu p^+ \rightarrow \phi_2(\bar{\nu}_e)^* p^+ \rightarrow \phi_2 e^+ n$. These explanations of MiniBooNE, however, are excluded by solar antineutrino searches [41] and meson decays [42]. We note that a decaying-sterile neutrino has been recently searched for by the IceCube Neutrino Observatory, finding a preference for decay [43]. Though the analysis uses invisible decay, it is

largely insensitive to the decay being visible or invisible due to the steeply falling spectra as discussed in [38]; see [44] for a recent discussion of the latter scenario.

C. Future prospects

This study focuses specifically on the possibility of a $\bar{\nu}_e$ excess in the BNB; thus, it is relevant to consider whether future experiments along the BNB will be sensitive to such an excess. MicroBooNE is part of the short baseline neutrino program at Fermilab, which includes the upcoming ICARUS and SBND experiments [45]. ICARUS and SBND also use LArTPC detectors, so they too will suffer from the $\bar{\nu}_e$ -Ar cross section suppression at low energies. Of the SBN experiments, SBND in particular is the most optimistic setup to search for a $\bar{\nu}_e$ excess. This is because it will benefit from a tenfold enhancement in event rate compared to MicroBooNE, as it is situated closer to the BNB target [45]. Even so, assuming SBND reconstruction efficiency is similar to MicroBooNE, a factor of ~ 10 enhancement in the backgrounds and excess templates shown in the right panels of Fig. 4 will not significantly improve the sensitivity to *Scenario 3*, unless a substantial reduction of the backgrounds can be achieved.

In view of the challenges in detecting antineutrinos in LArTPC detectors along the BNB, we turn to a different kind of detector for this measurement; the Accelerator Neutrino Neutron Interaction Experiment (ANNIE). ANNIE is a 26-ton water Cherenkov detector located at 100 m from the BNB target [46,47]. The water volume is followed by a muon-range detector to allow the detection of muon neutrino and antineutrino interactions. One of the primary goals of ANNIE is to measure the neutron multiplicity in CCQE interactions. For that, the detector is doped with gadolinium, so that neutrons produced in neutrino events can be detected via delayed ~ 8 MeV photons emitted in neutron-gadolinium capture.

While ANNIE has so far focused on muon events, the detector can also measure inverse beta-decay events, $\bar{\nu}_e p^+ \rightarrow e^+ n$. The signature is a single, low-energy positron Cherenkov ring with a delayed neutron capture. The photo coverage of ANNIE phase-II, in particular, provides the right environment for this measurement and can be used on a search for a $\bar{\nu}_e$ interpretation of the LEE at the current location of the detector. A detailed background study is needed to estimate ANNIE's sensitivity to the excess flux in Fig. 7. Nevertheless, the addition of water-based liquid scintillator in the detector volume, ANNIE phase-III, would be a clear improvement to mitigate backgrounds [47]. Finally, because of the sheer magnitude of the excess of $\bar{\nu}_e$ required by the MiniBooNE LEE and the lack of information on its size below $E_\nu < 200$ MeV, ANNIE can start to probe $\bar{\nu}_e$ -based explanations of the MiniBooNE LEE even if it is unable to detect the intrinsic $\bar{\nu}_e$ flux.

VI. CONCLUSION

The MicroBooNE experiment has not found any evidence for an electron-neutrino interpretation of the MiniBooNE low-energy excess (LEE). In this article, we show that this fact can be reconciled with the MiniBooNE observation if the LEE is caused by electron-antineutrinos instead. This is due to three main reasons: (i) two out of the three MicroBooNE analyses have focused on single proton final states, and these are rarely produced in antineutrino-nucleus scattering; (ii) the energy of the initial $\bar{\nu}_e$ is substantially under-reconstructed in MicroBooNE due to the invisible final state neutron; and (iii) the antineutrino cross sections per nucleon on ^{40}Ar are suppressed with respect to those in CH_2 , the nuclear targets in MiniBooNE. The differences in total cross section are due to the difference in the proton-to-neutron ratio (4:3 at MiniBooNE compared to 9:10 at MicroBooNE), but, more importantly, due to the larger size of the argon nucleus and its nonisoscalar nature. Contrary to carbon, the proton separation energy in argon is larger than that of neutrons, requiring a larger energy transfer in $\bar{\nu}$ CCQE scattering. As shown in the bottom panel of Fig. 2, this threshold effect is particularly significant in the energy region of the MiniBooNE LEE.

To quantify the impact of a $\bar{\nu}_e$ -interpretation of the LEE on the latest MicroBooNE results, we followed Ref. [17] and reproduced the results of the MicroBooNE Wire-Cell template analysis [13]. Because antineutrinos do not produce protons, we focused on the inclusive Wire-Cell sample, which does not require a proton connected to the neutrino interaction vertex. We then estimated new detector response matrices under the assumption of electron-antineutrino charged-current scattering and proceeded to unfold the MiniBooNE LEE into a $\bar{\nu}_e$ LEE template, shown in the top panel of Fig. 3. We checked that the unfolded template reproduces the MiniBooNE LEE once folded back into MiniBooNE with our response matrix, as shown in the bottom panel of Fig. 3.

We showed that if the antineutrino to neutrino ratio of the LEE event rate is 100% ($f_{\bar{\nu}/\nu} = 1$), then MicroBooNE's sensitivity is significantly reduced to less than 2σ (assuming Wilks' theorem), as much fewer LEE events are expected in the detector. If the number of antineutrino-induced LEE events is 50%, then MicroBooNE's sensitivity to the nominal template is reduced to less than 3σ CL for the nominal LEE template. As pointed out in Ref. [17], due to the large background systematic uncertainties in MiniBooNE, choosing different templates with an excellent fit to the LEE, $p_{\text{val}}^{\text{LEE}} > 90\%$, can have a significant impact on MicroBooNE's sensitivity. Using the best and worst-case template choices from Ref. [17], if we take the MiniBooNE excess at face value and require $x = 1.0$ the MicroBooNE data constrain the antineutrino-to-neutrino fraction of the LEE event rate to be at least 0.72 for template T1 (best-case scenario) and 0.24 for template T2

(worst-case scenario) at the 2σ CL, assuming Wilks' theorem for 1 d.o.f.

The Deep-Learning analysis focused specifically on the $1e1p$ event topology, maximizing its sensitivity to electron-neutrino CCQE events. While this analysis had the largest purity, the requirement of a final state proton makes it insensitive to antineutrinos, which can only produce a proton through nuclear final state interactions. Finally, the Pandora analysis focused on pionless topologies, separating them into $1e0\pi Np$ and $1e0\pi 0p$. While the latter does not require a proton in the final state, it is also the least sensitive and less pure of the analyses. It is also the only one that observes a small excess of events. We can then conclude that the choice of event topologies makes the Pandora and Deep-Learning analyses insensitive to an excess of $\bar{\nu}_e$.

The other LArTPC detectors along the BNB, SBND, and ICARUS, will face the same issues as MicroBooNE when testing a $\bar{\nu}_e$ explanation of the MiniBooNE LEE, due to the $\bar{\nu}_e$ -Ar cross section suppression at low energy. In principle, the near detector of the SBN program, SBND, will observe more antineutrino events but on top of a higher overall event rate. In light of this, we have pointed out a different possibility to directly search for $\bar{\nu}_e$ in the BNB using the ANNIE detector. The phase-II of ANNIE is particularly well-suited for the study of low-energy inverse beta decay. The scattering of antineutrinos on free protons inside the water-based Cherenkov detector produces a prompt positron signal followed by a delayed capture of neutrons on gadolinium. A detailed study of the backgrounds is needed to assess the final sensitivity of the experiment. Further improvements would be possible with phase III, where the separation of scintillation and Cherenkov light could be achieved with a water-based liquid-scintillator volume. A dedicated analysis at ANNIE can shed new light on SM as well as BSM explanations of the MiniBooNE LEE, targeting the BNB flux at energies of the LEE and below.

ACKNOWLEDGMENTS

We want to thank Mayly C. Sanchez and Michael Wurm for interesting discussions on the ANNIE experiment. The research at the Perimeter Institute is supported in part by the Government of Canada through NSERC and by the Province of Ontario through the Ministry of Economic Development, Job Creation and Trade, MEDT. C. A. A. is supported by the Faculty of Arts and Sciences of Harvard University. N. W. K. is supported by the NSF Graduate Research Fellowship under Grant No. 1745302. J. M. C. and M. H. S. are supported by the NSF.

APPENDIX: FURTHER DETAILS ON THE ANALYSIS

This supplement to the main text is intended to provide additional detail on the analysis, specifically concerning the

calculation of the response matrices used to unfold the MiniBooNE excess and obtain a $\bar{\nu}_e$ prediction in MicroBooNE.

1. Response matrix calculation

To approximate $R_{i\alpha}$ in MiniBooNE, we first calculate the conditional probability density function (PDF) $P(E_{\bar{\nu}}^{\text{reco}}|E_{\bar{\nu}}^{\text{true}})$. We marginalize over $E_{\nu}^{\text{QE,true}}$, which, due

$$\begin{aligned} P(E_{\bar{\nu}}^{\text{reco}}|E_{\bar{\nu}}^{\text{true}}) &= \int_0^\infty dE_{\nu}^{\text{QE,true}} P(E_{\nu}^{\text{QE,reco}}|E_{\nu}^{\text{QE,true}}) P(E_{\nu}^{\text{QE,true}}|E_{\bar{\nu}}^{\text{true}}) \\ &= \sum_{E_{\nu}^{\text{QE,true}} \text{ bins } k} P_k^{\text{GENIE}}(E_{\bar{\nu}}^{\text{true}}) \int_{(E_{\nu}^{\text{QE,true}})_k^{\text{low}}}^{(E_{\nu}^{\text{QE,true}})_k^{\text{high}}} dE_{\nu}^{\text{QE,true}} \frac{\exp\left[-\frac{(E_{\nu}^{\text{QE,reco}} - E_{\nu}^{\text{QE,true}})^2}{2(\sigma(E_{\bar{\nu}}^{\text{true}}))^2}\right]}{\sqrt{2\pi}(\sigma(E_{\bar{\nu}}^{\text{true}}))^2} \\ &= \sum_{E_{\nu}^{\text{QE,true}} \text{ bins } k} P_k^{\text{GENIE}}(E_{\bar{\nu}}^{\text{true}}) \frac{1}{2} \left[\text{Erf}\left(\frac{(E_{\nu}^{\text{QE,true}})_k^{\text{high}} - E_{\nu}^{\text{QE,reco}}}{\sqrt{2}\sigma(E_{\bar{\nu}}^{\text{true}})}\right) - \text{Erf}\left(\frac{(E_{\nu}^{\text{QE,true}})_k^{\text{low}} - E_{\nu}^{\text{QE,reco}}}{\sqrt{2}\sigma(E_{\bar{\nu}}^{\text{true}})}\right) \right], \quad (\text{A1}) \end{aligned}$$

where Erf denotes the error function and we define $P_k^{\text{GENIE}}(E_{\bar{\nu}}^{\text{true}}) \equiv P((E_{\nu}^{\text{QE,true}})_k|E_{\bar{\nu}}^{\text{true}})$ to be the binned $E_{\nu}^{\text{QE,true}}$ probability distribution calculated via GENIE v3.02.00 for a given $E_{\bar{\nu}}^{\text{true}}$. In the third line, we convert the continuous integral over $E_{\nu}^{\text{QE,true}}$ to a discrete sum over the binned P_k^{GENIE} distribution and explicitly show our Gaussian approximation, where we consider a flat 17% energy

to nuclear effects and differences between the E_{ν}^{QE} and $E_{\bar{\nu}}^{\text{QE}}$ expressions, is not necessarily the same as the generated antineutrino energy $E_{\bar{\nu}}^{\text{true}}$. We use GENIE v3.02.00 to discretely approximate $P(E_{\nu}^{\text{QE,true}}|E_{\bar{\nu}}^{\text{true}})$ and use a Gaussian approximation for $P(E_{\nu}^{\text{QE,reco}}|E_{\nu}^{\text{QE,true}})$. Note that since we are considering MiniBooNE neutrino mode data, $E_{\bar{\nu}}^{\text{reco}} \equiv E_{\nu}^{\text{QE,reco}}$. The full calculation of the conditional PDF is

resolution, i.e., $\sigma(E_{\bar{\nu}}^{\text{true}}) = 0.17E_{\bar{\nu}}^{\text{true}}$. We also consider a flat 5% $E_{\bar{\nu}}^{\text{QE}}$ underestimation bias. This energy resolution and bias are derived from the most recent MiniBooNE ν_e data release [48]; for more details see Appendix A 2. In the fourth line, we integrate the Gaussian within each $E_{\nu}^{\text{QE,true}}$ bin k . We can then approximate $R_{i\alpha}$ by integrating $P(E_{\bar{\nu}}^{\text{reco}}|E_{\bar{\nu}}^{\text{true}})$ over $E_{\bar{\nu}}^{\text{reco}}$ bin i and averaging over $E_{\bar{\nu}}^{\text{true}}$ bin α ,

$$R_{i\alpha} = \frac{1}{(E_{\bar{\nu}}^{\text{true}})_\alpha^{\text{high}} - (E_{\bar{\nu}}^{\text{true}})_\alpha^{\text{low}}} \int_{(E_{\bar{\nu}}^{\text{true}})_\alpha^{\text{low}}}^{(E_{\bar{\nu}}^{\text{true}})_\alpha^{\text{high}}} \int_{(E_{\bar{\nu}}^{\text{reco}})_i^{\text{low}}}^{(E_{\bar{\nu}}^{\text{reco}})_i^{\text{high}}} dE_{\bar{\nu}}^{\text{reco}} dE_{\bar{\nu}}^{\text{true}} P(E_{\bar{\nu}}^{\text{reco}}|E_{\bar{\nu}}^{\text{true}}). \quad (\text{A2})$$

The $\bar{\nu}_e$ response matrix in MicroBooNE begins with a calculation similar to Eq. (A1). Due to the calorimetric energy reconstruction used in the Wire-Cell analysis, the calculation of the conditional PDF $P(E_{\bar{\nu}}^{\text{reco}}|E_{\bar{\nu}}^{\text{true}})$ in MicroBooNE marginalizes over the true positron kinetic energy,

$$P(E_{\bar{\nu}}^{\text{reco}}|E_{\bar{\nu}}^{\text{true}}) = \int_0^\infty dT_{e^+}^{\text{true}} P(E_{\bar{\nu}}^{\text{Cal,reco}}|T_{e^+}^{\text{true}}) P(T_{e^+}^{\text{true}}|E_{\bar{\nu}}^{\text{true}}). \quad (\text{A3})$$

We perform the same procedure used for MiniBooNE to approximate $R_{i\alpha}$ in MicroBooNE, leveraging GENIE v3.02.00 to calculate $P(T_{e^+}^{\text{true}}|E_{\bar{\nu}}^{\text{true}})$ and making a Gaussian approximation for $P(E_{\bar{\nu}}^{\text{Cal,reco}}|T_{e^+}^{\text{true}})$. Following Ref. [16], we consider a flat 2% relative bias and 12% relative uncertainty on the lepton kinetic energy. The MiniBooNE and MicroBooNE $\bar{\nu}_e$ response matrices calculated using Eqs. (A1)–(A3) are shown in Fig. 8. The significant

population of events below the $E_{\bar{\nu}}^{\text{reco}} = E_{\bar{\nu}}^{\text{true}}$ line in the MicroBooNE matrix indicates the underestimation bias from the invisible neutron.

2. Approximations in response matrices

To approximate the truth-level PDFs in Eqs. (A1) and (A3), we use the latest version of the GENIE event generator, GENIE v3.02.00 [18]. As the hypotheses presented in this study attribute at least part of the MiniBooNE excess to CC $\bar{\nu}_e$ scattering, we generate all charged-current $\bar{\nu}_e$ interactions included in GENIE v3.02.00 in both CH₂ and Ar40. Truth-level distributions of final state variables in these interactions are shown in Fig. 9, including the positron kinetic energy T_{e^+} , neutron kinetic energy T_n , and positron scattering angle $\cos\theta_{e^+}$. We also show the truth-level distribution of the energy reconstruction definition for each experiment— $E_{\bar{\nu}}^{\text{QE}}$ for MiniBooNE and $E_{\bar{\nu}}^{\text{Cal}}$ for MicroBooNE. Note that in the MicroBooNE case, the $E_{\bar{\nu}}^{\text{Cal}}$ distribution matches the T_{e^+} distribution, as the final state




FIG. 8. MiniBooNE (top) and MicroBooNE (bottom) $\bar{\nu}_e$ response matrices R_{ia} , calculated according to Appendix A 1. Reconstructed energy binning in each matrix reflects the binning reported by each collaboration.

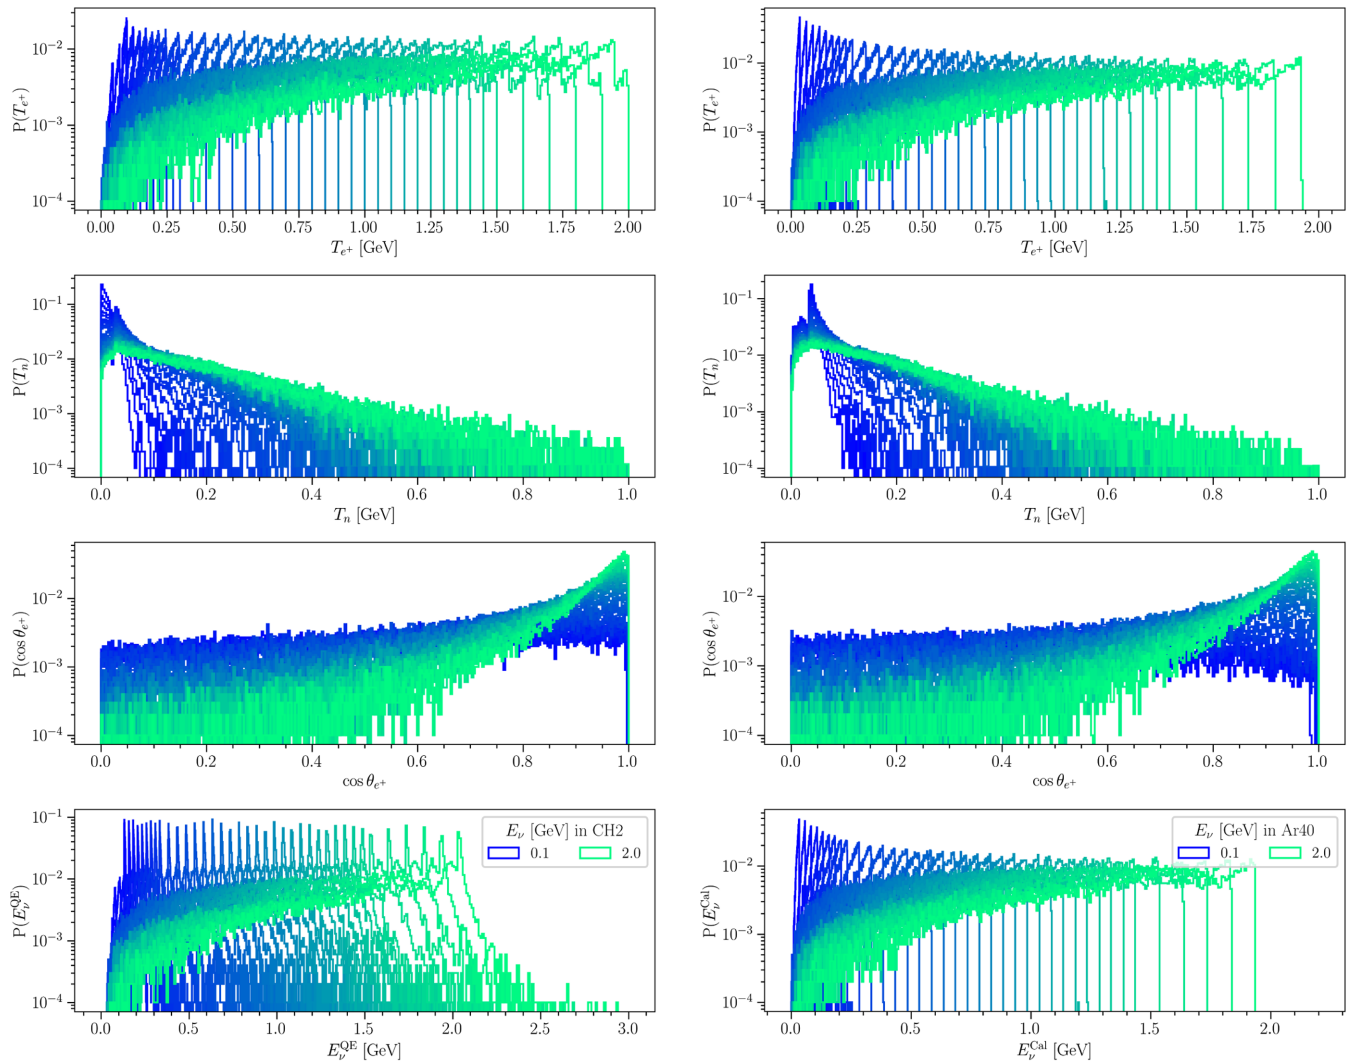


FIG. 9. Truth-level distributions of final state kinematic variables for $\bar{\nu}_e$ scattering in CH_2 (left) and Ar40 (right), generated using GENIE v3.02.00.

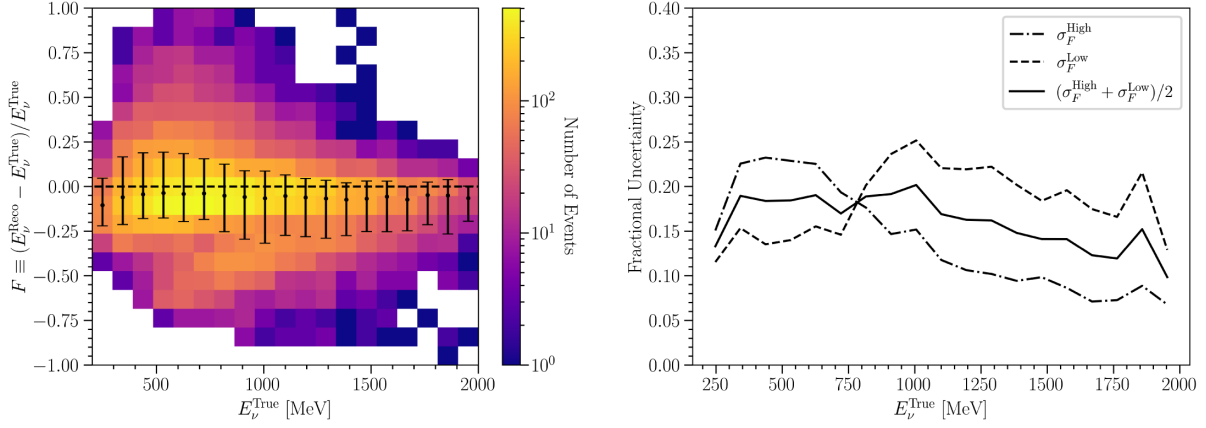


FIG. 10. Left: 2D distribution of the fractional energy reconstruction error F as a function of the true neutrino energy. The black points and error bars indicate the median and $\pm 1\sigma$ extent, respectively, of the F distribution in each E_ν^{True} bin. Right: upper and lower uncertainties on F as a function of the true neutrino energy. The average of the two is also shown.

neutron is invisible. In principle, final state protons and charged pions generated in $\bar{\nu}_e$ scattering via interactions within nuclear medium will also contribute to E_ν^{Cal} ; however, we have ignored this effect for the purposes of this study.

The MiniBooNE energy resolution is approximated using the latest $\nu_\mu \rightarrow \nu_e$ Monte Carlo data release from the MiniBooNE Collaboration [48]. In Fig. 10, we show a 2D histogram of the fractional error of each event,

$$F \equiv \frac{E_\nu^{\text{Reco}} - E_\nu^{\text{True}}}{E_\nu^{\text{True}}}, \quad (\text{A4})$$

as a function of the true neutrino energy. A profiled version of this distribution is overlaid in black, where the data points and error bars indicate the median and $\pm 1\sigma$ extent, respectively, of the F distribution in each E_ν^{True} bin. From the median, one can see that MiniBooNE tends to consistently underpredict the ν_e energy. The fractional uncertainty on the neutrino energy is also relatively constant across the relevant energy range, though it is not symmetric. This is shown in the right panel of Fig. 10, which plots the upper and lower 1σ fractional uncertainty on the neutrino energy as a function of E_ν^{True} , as well as the average of the two. It is apparent that the uncertainty on F is relatively flat across the relevant true neutrino energy range. One can see that for $E_\nu^{\text{True}} < 750$ MeV, the F distribution has a larger extend above the median, while the opposite is true for $E_\nu^{\text{True}} > 750$ MeV. While this is an interesting effect, for the purpose of this study we approximate the MiniBooNE E_ν fractional uncertainty to be 17%—the average value σ_F across the full neutrino energy range. We also incorporate a 5% underprediction bias, which is the median value of F across the full neutrino energy range. In principle, $\bar{\nu}_e$ scattering will behave differently than ν_e scattering in MiniBooNE, as the momentum transfer Q^2 distribution and final state lepton kinematics differ between

the two. This effect is accounted for in the construction of the response matrices for both MiniBooNE and MicroBooNE—however, we ignore it when approximating the MiniBooNE energy resolution, as it is a subdominant effect here. The fractional uncertainty on positron EM showers in MicroBooNE is taken to be 12% with a 2% underprediction bias, as quoted in Ref. [16].

As mentioned in the main text, we use the provided electron reconstruction efficiency $\epsilon(E_{e^+})$ in MiniBooNE [21] to approximate the $\bar{\nu}_e$ reconstruction efficiency ϵ_α . As MiniBooNE is a spherically symmetric detector that can only reconstruct the final state lepton, the e^+ energy is the dominant effect in the (anti)neutrino detection efficiency. Using GENIE v3.02.00, we can estimate $P(E_{e^+}|E_{\bar{\nu}})$ as shown in Fig. 9 and thus approximate the lepton energy-averaged detection efficiency,

$$\epsilon(E_{\bar{\nu}}) = \int_0^\infty \epsilon(E_{e^+})P(E_{e^+}|E_{\bar{\nu}})dE_{e^+}. \quad (\text{A5})$$




FIG. 11. MiniBooNE $\bar{\nu}_e$ detection efficiency as a function of the true antineutrino energy.

This efficiency is shown in Fig. 11. For reasons discussed in the main text, we consider the $\bar{\nu}_e$ detection efficiency in MicroBooNE to be the same as the ν_e detection efficiency released by the Collaboration. This likely overestimates the $\bar{\nu}_e$ rate in MicroBooNE, as the neutron created in $\bar{\nu}_e$

CC interactions is not visible in the detector. This would make our estimation of MicroBooNE's sensitivity to a $\bar{\nu}_e$ -generated excess artificially strong. Thus, it is a conservative assumption, given the conclusion of this paper—that MicroBooNE is not sensitive to a $\bar{\nu}_e$ -generated excess.

-
- [1] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Phys. Rev. D* **79**, 072002 (2009).
- [2] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Phys. Rev. D* **103**, 052002 (2021).
- [3] V. Brdar and J. Kopp, *Phys. Rev. D* **105**, 115024 (2022).
- [4] K. J. Kelly and J. Kopp, [arXiv:2210.08021](https://arxiv.org/abs/2210.08021).
- [5] M. A. Acero *et al.*, [arXiv:2203.07323](https://arxiv.org/abs/2203.07323).
- [6] A. Aguilar-Arevalo *et al.* (LSND Collaboration), *Phys. Rev. D* **64**, 112007 (2001).
- [7] H. Chen *et al.* (MicroBooNE Collaboration), Proposal for a new experiment using the Booster and NuMI Neutrino beamlines: MicroBooNE (2007), <https://inspirehep.net/literature/776376>.
- [8] P. Abratenko *et al.* (MicroBooNE), *J. Instrum.* **16**, P02017 (2021).
- [9] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. Lett.* **128**, 241801 (2022).
- [10] P. Abratenko *et al.* (MicroBooNE Collaboration), 10.2172/1573217.
- [11] G. D'Agostini, *Nucl. Instrum. Methods Phys. Res., Sect. A* **362**, 487 (1995).
- [12] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **105**, 112003 (2022).
- [13] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **105**, 112005 (2022).
- [14] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **103**, 052012 (2021).
- [15] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **103**, 092003 (2021).
- [16] P. Abratenko *et al.* (MicroBooNE Collaboration), *J. Instrum.* **17**, P01037 (2022).
- [17] C. A. Argüelles, I. Esteban, M. Hostert, K. J. Kelly, J. Kopp, P. A. N. Machado, I. Martinez-Soler, and Y. F. Perez-Gonzalez, *Phys. Rev. Lett.* **128**, 241802 (2022).
- [18] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, [arXiv:1510.05494](https://arxiv.org/abs/1510.05494).
- [19] L. Alvarez-Ruso *et al.* (GENIE Collaboration), *Eur. Phys. J. Spec. Top.* **230**, 4449 (2021).
- [20] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **105**, 072001 (2022).
- [21] Z. Pavlovic, R. Van de Water, and S. Zeller, MiniBooNE gamma-ray and electron efficiencies, Technical Report, 2012.
- [22] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Phys. Rev. Lett.* **121**, 221801 (2018).
- [23] S. S. Wilks, *Ann. Math. Stat.* **9**, 60 (1938).
- [24] B. Pontecorvo, *Zh. Eksp. Teor. Fiz.* **53**, 1717 (1967).
- [25] A. Cisneros, *Astrophys. Space Sci.* **10**, 87 (1971).
- [26] M. B. Voloshin and M. I. Vysotsky, *Sov. J. Nucl. Phys.* **44**, 544 (1986).
- [27] L. B. Okun, *Sov. J. Nucl. Phys.* **44**, 546 (1986).
- [28] M. Agostini *et al.* (Borexino Collaboration), *Astropart. Phys.* **125**, 102509 (2021).
- [29] K. Abe *et al.* (Super-Kamiokande Collaboration), *Astropart. Phys.* **139**, 102702 (2022).
- [30] S. Abe *et al.* (KamLAND Collaboration), *Astrophys. J.* **925**, 14 (2022).
- [31] A. M. Cooper-Sarkar, J. G. Guy, A. G. Michette, M. Tyndel, and W. Venus, *Phys. Lett.* **112B**, 97 (1982).
- [32] V. Cirigliano and I. Rosell, *Phys. Rev. Lett.* **99**, 231801 (2007).
- [33] D. A. Bryman, M. S. Dixit, R. Dubois, J. A. Macdonald, T. Numao, B. Olaniyi, A. Olin, and J. M. Poutissou, *Phys. Rev. D* **33**, 1211 (1986).
- [34] G. Czapek *et al.*, *Phys. Rev. Lett.* **70**, 17 (1993).
- [35] D. I. Britton *et al.*, *Phys. Rev. D* **49**, 28 (1994).
- [36] A. Aguilar-Arevalo *et al.* (PiENU Collaboration), *Phys. Rev. Lett.* **115**, 071801 (2015).
- [37] S. Palomares-Ruiz, S. Pascoli, and T. Schwetz, *J. High Energy Phys.* **09** (2005) 048.
- [38] Z. Moss, M. H. Moulai, C. A. Argüelles, and J. M. Conrad, *Phys. Rev. D* **97**, 055017 (2018).
- [39] M. Dentler, I. Esteban, J. Kopp, and P. Machado, *Phys. Rev. D* **101**, 115013 (2020).
- [40] A. de Gouvêa, O. L. G. Peres, S. Prakash, and G. V. Stenico, *J. High Energy Phys.* **07** (2020) 141.
- [41] M. Hostert and M. Pospelov, *Phys. Rev. D* **104**, 055031 (2021).
- [42] J. M. Berryman, A. De Gouvêa, K. J. Kelly, and Y. Zhang, *Phys. Rev. D* **97**, 075030 (2018).
- [43] R. Abbasi *et al.* (IceCube Collaboration), *Phys. Rev. Lett.* **129**, 151801 (2022).
- [44] J. M. Hardin, I. Martinez-Soler, A. Diaz, M. Jin, M. W. Kamp, C. A. Argüelles, J. M. Conrad, and M. H. Shaevitz, [arXiv:2211.02610](https://arxiv.org/abs/2211.02610).
- [45] M. Antonello *et al.* (MicroBooNE, LAr1-ND, ICARUS-WA104 Collaborations), [arXiv:1503.01520](https://arxiv.org/abs/1503.01520).
- [46] I. Anghel *et al.* (ANNIE Collaboration), [arXiv:1504.01480](https://arxiv.org/abs/1504.01480).
- [47] A. R. Back *et al.* (ANNIE Collaboration), [arXiv:1707.08222](https://arxiv.org/abs/1707.08222).
- [48] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), [arXiv:2110.15055](https://arxiv.org/abs/2110.15055).

Chapter 6

Neutrissimos: Heavy Neutral Leptons with a Dipole Moment

In this chapter, we discuss a phenomenological explanation of the MiniBooNE excess involving both an eV-scale sterile neutrino and an MeV-scale heavy neutral lepton (HNL) with an effective transition magnetic moment coupling to active neutrinos. This “mixed model” was introduced in Ref. [27] and refined in Ref. [31]. The HNL in this model is hereafter referred to as a “neutrissimo” [31], a term which has previously been used to describe interacting HNLs in the literature [215, 216].

We begin with a brief overview of the theoretical framework behind neutrissimos. We then motivate the mixed model and discuss its implications for global neutrino experiments. A robust evaluation of the neutrissimo signal in MiniBooNE is performed. For a neutrissimo mass $m_{\mathcal{N}} \sim 500$ MeV, single photons from visible neutrissimo decay are shown to provide a reasonable explanation of the energy and angular distributions of the MiniBooNE excess. World-leading constraints on neutrissimos are derived using elastic scattering data from the MINER ν A experiment [217–219]. While the MINER ν A data exclude a large region of parameter space for $m_{\mathcal{N}} = \mathcal{O}(100)$ MeV, they do not rule out the MiniBooNE solution at the 95% confidence level. The details behind this analysis are given in Ref. [31]; the full published manuscript is included in section 6.4.

Publications covered in this chapter for which I held a leading role: [27, 31]

6.1 Dipole-Portal Neutrissimos

The neutrissimo \mathcal{N} is a right-handed neutrino that couples to the active neutrinos through a transition magnetic moment. This coupling is described by the dimension-five dipole operator

$$\mathcal{L}_D = d_{\alpha\mathcal{N}} \bar{\nu}_\alpha \sigma_{\mu\nu} F^{\mu\nu} \mathcal{N}_R + \text{h.c.}, \quad (6.1)$$

where $d_{\alpha\mathcal{N}}$ is the dipole coupling between the right-handed neutrissimo \mathcal{N}_R and the left-handed weak flavor eigenstate ν_α , and $F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu$ is the EM field strength tensor. The dipole coupling has units of inverse energy, as equation (6.1) is an effective operator. This means it is agnostic to the ultraviolet (UV) completion: the short-distance physics generating the dipole coupling at a higher energy scale. Possible UV completions for the dipole operator include (but are not limited to) $SU(2)_L \times SU(2)_R$ theories [220], leptoquark models [81], and models introducing additional Higgs multiplets or other charged scalars [221–224]. We will treat neutrissimos within the effective field theory framework throughout this chapter. Figure 6-1 shows a Feynman diagram of the effective operator in equation (6.1). It facilitates a number of new interactions involving the neutrissimo, the most relevant of which are shown in figure 6-2.

Equation (6.1) is only valid below the electroweak scale, as it does not respect $SU(2)_L \times U(1)_Y$ gauge invariance. Before electroweak symmetry breaking (EWSB), the dipole coupling is described by the gauge invariant dimension six operator

$$\mathcal{L} \supset \frac{1}{\Lambda^2} \bar{L}_\alpha \tilde{H} \sigma^{\mu\nu} \mathcal{N}_R (C_B^\alpha B_{\mu\nu} + C_W^\alpha W_{\mu\nu}^a \sigma_a) + \text{h.c.}, \quad (6.2)$$

where Λ is the UV completion scale, L_α is the left-handed lepton $SU(2)_L$ doublet, $\tilde{H} = i\sigma_2 H^*$ is the conjugate Higgs field, and C_B^α and C_W^α are the Wilson coefficients for couplings to the B_μ and W_μ vector gauge bosons, respectively. The dipole coupling in equation (6.1) after EWSB is then given by

$$d_{\alpha\mathcal{N}} = \frac{v_h}{\sqrt{2}\Lambda^2} (C_B^\alpha \cos \theta_W + C_W^\alpha \sin \theta_W), \quad (6.3)$$

where v_h is the Higgs vacuum expectation value and θ_W is the Weinberg angle. Note that couplings are also generated with the W and Z bosons after EWSB, but interactions involving these couplings are suppressed by G_F and are thus not as important as the photon coupling.

Transition magnetic moments between different neutrino species have been discussed extensively in the literature [81, 130, 223, 225–230]. The imprint of heavy neutral leptons with transition magnetic moments in experimental data has received particular attention [81, 123, 124, 126, 127, 130, 135, 231–240]. One potential complication of this model is that large transition magnetic moments often lead to large Dirac masses $m_{\nu\mathcal{N}}$ for neutrinos [31, 81, 130]. This is because the photon line can typically be removed from figure 6-1, resulting in a contribution to the Dirac mass term $\mathcal{L} \supset m_{\nu\mathcal{N}} \bar{\nu}_L \mathcal{N}_R$. Dipole coupling strengths of interest to MiniBooNE ($d \sim 10^{-6}$ GeV) would generate Dirac mass contributions $m_{\nu\mathcal{N}} = \mathcal{O}(\text{MeV})$ [81]. Considering a neutrino mass $m_{\mathcal{N}} \sim 500$ MeV, the typical seesaw relation of equation (1.23) would predict active neutrino masses $m_\nu \gg 1$ eV, in conflict with existing limits. This can be remedied by the inverse seesaw mechanism [241, 242], in which the smallness of active neutrino masses comes from approximate lepton number conservation without relying on large right-handed neutrino masses. [31, 81].

Another complication is the potential existence of mass mixing between the neutrino and the active neutrinos. This mixing is heuristically given by $U_{\alpha\mathcal{N}} \sim m_{\nu\mathcal{N}}/m_{\mathcal{N}}$ and is large even in the inverse seesaw mechanism [81]. If this operator is not suppressed, the interactions of the neutrino are dominated by those involving the W and Z bosons. Any potential mass mixing would be subject to constraints from accelerator neutrino experiments, which can be strong for $m_{\mathcal{N}} = \mathcal{O}(100 \text{ MeV})$ [243]. Specifically, searches for HNLs produced in kaon decay at T2K (using the ND280 near detector) [244], E949 [245], and PS191 [246] set upper bounds on the muon-flavor mixing of $|U_{\mu\mathcal{N}}|^2 \leq 10^{-8} - 10^{-9}$ up to the kaon mass threshold of $m_{\mathcal{N}} \sim 400$ MeV. Within the range $400 \leq m_{\mathcal{N}} [\text{MeV}] \leq 2000$, NuTeV is able to set upper bounds of $|U_{\mu\mathcal{N}}|^2 \leq 10^{-6} - 10^{-7}$ by searching for HNLs produced in D meson decays.

In order to suppress mass mixing, the Dirac mass itself must be small compared




Figure 6-1: Feynman diagram depicting the effective dipole operator of equation (6.1).

to the neutrino mass. One possible method to accomplish this is the approximate horizontal $SU(2)_H$ symmetry proposed by Voloshin [247]. The idea of an approximate $SU(2)_H$ symmetry initially gained popularity as a means to enable the ν_e diagonal/transition magnetic moment solution to the solar neutrino problem [248–250]. For the heavy neutral lepton case, we consider a scenario in which (ν_L^C, \mathcal{N}_R) transforms as a doublet under $SU(2)_H$ [81, 224, 251]. The antisymmetric nature of the dipole operator $\bar{\nu}_\alpha \sigma_{\mu\nu} F^{\mu\nu} \mathcal{N}_R$ in flavor space means that it transforms like a singlet under $SU(2)_H$, while the symmetric nature of the Dirac mass term makes it a triplet under $SU(2)_H$. Thus, approximate $SU(2)_H$ conservation will suppress the Dirac mass contribution from the photon-less version of figure 6-1 while allowing for large contributions to the transition magnetic moment. Though it is not clear the level of fine-tuning required within this model to accommodate the MiniBooNE solution [224], we consider this sufficient motivation to consider a scenario in which the mass mixing operator of the heavy neutrino is negligible compared to the dipole operator. It is also worth pointing out the leptoquark UV completion of the dipole operator proposed in Ref. [81], which can naturally suppress the mass mixing of neutrinos via an $SU(2)_H$ symmetry without the need for significant fine-tuning.

6.2 Overview of the Mixed Model

The mixed model consists of the eV-scale sterile neutrino described in section 1.4 and the MeV-scale neutrino described above in section 6.1. This model is motivated by




Figure 6-2: New interactions involving the neutrissimo that are enabled by the dipole operator in equation (6.1), including three-body π^0 decay (figure 6-2a), Primakoff-like upscattering (figure 6-2b), and neutrissimo decay (figure 6-2c).

the significant tension observed between appearance and disappearance experiments in $3 + 1$ global fits [11, 12, 93, 252], which comes from strong signals indicating ν_e appearance and ν_e disappearance coupled with the lack of such a signal in the ν_μ disappearance channel. This tension is measured using the parameter goodness-of-fit (PG) test [253], which uses the test statistic

$$\begin{aligned}\chi_{PG}^2 &= \chi_{\text{global}}^2 - (\chi_{\text{app}}^2 + \chi_{\text{dis}}^2), \\ N_{PG} &= (N_{\text{app}} + N_{\text{dis}}) - N_{\text{global}},\end{aligned}\tag{6.4}$$

where χ_{global}^2 , χ_{app}^2 , and χ_{dis}^2 are the total χ^2 at the $3 + 1$ best-fit parameters for global, appearance-only, and disappearance-only experiments, respectively. The number of degrees-of-freedom N is defined similarly. The probability that the appearance and disappearance experiments come from the same underlying $3 + 1$ model can be calculated using the χ_{PG}^2 distribution for N_{PG} degrees-of-freedom.

This tension appears to be driven significantly by MiniBooNE data [12, 27]. Within the Harvard-Columbia-MIT global fit that existed at the time of Ref. [27], the PG test tension decreased from 4.8σ to 2.5σ upon removing the MiniBooNE dataset. The allowed regions in Δm_{41}^2 - $\sin^2 2\theta_{\mu e}$ parameter space for the global, appearance-only, and disappearance-only experiments both before and after removing MiniBooNE are shown in figure 6-3 and figure 6-4, respectively. One can see visually that there is more overlap in the Δm^2 dimension between the three different allowed regions after removing MiniBooNE. This is likely because the MiniBooNE excess peaks strongly

at low energies, an effect which is difficult to produce in the $3 + 1$ model, as discussed in section 2.2. The latest MiniBooNE $3 + 1$ fits discussed in section 5.3.1 specifically show that the oscillation hypothesis under-predicts the excess in the lowest bins of the MiniBooNE electron-like dataset. This behavior drives the global fit to the $\Delta m_{41}^2 < 1 \text{ eV}^2$ region, in contrast to the higher values ($\Delta m_{41}^2 \gtrsim 1 \text{ eV}^2$) preferred by the reactor and gallium anomalies.

The global fit without MiniBooNE, shown in the leftmost panel of figure 6-4, strongly prefers a solution at $\Delta m_{41}^2 \sim 1 \text{ eV}^2$ and $10^{-4} \lesssim \sin^2 2\theta_{\mu e} \lesssim 10^{-3}$. The best fit point is specifically at $\Delta m_{41}^2 = 1.3 \text{ eV}^2$ and $\sin^2 2\theta_{\mu e} = 6.9 \times 10^{-4}$. We take these as our benchmark $3 + 1$ parameters, which fix the $\nu_\mu \rightarrow \nu_e$ oscillation contribution to the MiniBooNE electron-like channel. The oscillation prediction alone is not sufficient to explain the MiniBooNE LEE; it can only accommodate $\sim 10\%$ of the total excess events and peaks at $E_\nu \sim 600 \text{ MeV}$, inconsistent with the low-energy nature of the MiniBooNE excess. However, the benefit of these specific oscillation parameters is that they can explain LSND, reactor, and gallium data while remaining consistent with null results from ν_μ disappearance searches. The impact of different mixing angles within the allowed region of the global fit without MiniBooNE is discussed in Ref. [27], and the impact of the (larger) oscillation contribution predicted by the MiniBooNE-MicroBooNE combined best fit [102] is discussed in Ref. [31].

The remaining majority of the MiniBooNE excess must then come from another source. In the mixed model, we ascribe the bulk of the LEE to the visible decays of neutrissimos. This model works well for MiniBooNE, as neutrissimos can be produced abundantly in the BNB through the first two diagrams in figure 6-2: Primakoff upscattering $\nu A \rightarrow \mathcal{N}A$ and the three-body decay $\pi^0 \rightarrow \gamma\nu\mathcal{N}$. The third diagram, neutrissimo decay $\mathcal{N} \rightarrow \nu\gamma$, could then provide a source of extra photons in MiniBooNE. As discussed in chapter 2, these photons would show up in the electron-like sample due to MiniBooNE’s reliance on Cherenkov rings for particle identification. Models of heavy neutrino decay to single photons have been previously considered as a solution to the MiniBooNE anomaly [123–135, 254]. We extend upon these studies by performing a robust analysis of the neutrissimo model in MiniBooNE within the




Figure 6-3: 3 + 1 global fits including MiniBooNE, considering global (left), appearance-only (middle), and disappearance-only (right) experiments. The allowed regions in 3 + 1 parameter space at the 90%, 95%, and 99% confidence levels are shown by the red, green, and blue points, respectively. The best-fit point is indicated by the star.




Figure 6-4: 3+1 global fits without MiniBooNE, considering global (left), appearance-only (middle), and disappearance-only (right) experiments. The allowed regions in 3 + 1 parameter space at the 90%, 95%, and 99% confidence levels are shown by the red, green, and blue points, respectively. The best-fit point is indicated by the star.

context of the mixed model. This analysis is described in detail in the next section. It turns out that compared to the 3+1 model, neutrissimo decays are more consistent with the energy and angular distributions of the MiniBooNE excess. Thus there are two distinct advantages of the mixed model: (1) it can provide an explanation for other anomalies in the neutrino sector while relieving tension in global 3 + 1 fits, and (2) it can provide a better explanation of the MiniBooNE excess itself.

6.3 Neutrissimos in MiniBooNE

The analysis presented in Ref. [27] used a custom simulation to estimate the single photon event rate from neutrissimo decays in the MiniBooNE detector. The updated

MiniBooNE analysis in Ref. [31] relied instead on the public `LeptonInjector`¹ software package, which will be described in more detail in section 6.3.1. In Ref. [27] we considered neutrino production via both $\pi^0 \rightarrow \gamma\nu\mathcal{N}$ decays and $\nu A \rightarrow \mathcal{N}A$ Primakoff upscattering. The former process was estimated using the Sanford-Wang parameterization of the π^0 production rate at the BNB [14,110] and was found to be negligible for the dipole parameters of interest to MiniBooNE. We therefore explicitly ignored this process, considering Primakoff upscattering as the sole neutrino production mechanism.

The $\nu_\alpha A \rightarrow \mathcal{N}A$ upscattering rate was calculated using the differential cross section with respect to the Mandelstam variable $t = Q^2 = (p_{\mathcal{N}} - p_\nu)^2$ [27,130,232,255],

$$\begin{aligned} \frac{d\sigma_{\nu_\alpha A \rightarrow \mathcal{N}A}}{dt} = \frac{2\alpha d_\alpha^2}{M} & \left[F_1^2(t) \left(\frac{1}{E_r} - \frac{1}{E_\nu} + m_{\mathcal{N}}^2 \frac{E_r - 2E_\nu - M}{4E_\nu^2 E_r M} + m_{\mathcal{N}}^4 \frac{E_r - M}{8E_\nu^2 E_r^2 M^2} \right) \right. \\ & \left. + \frac{F_2^2(t)}{4M^2} \left(\frac{2M}{E_\nu^2} ((2E_\nu - E_r)^2 - 2E_r M) + m_{\mathcal{N}}^2 \frac{E_r - 4E_\nu}{E_\nu^2} + \frac{m_{\mathcal{N}}^4}{E_\nu^2 E_r} \right) \right], \end{aligned} \quad (6.5)$$

where α is the fine structure constant, d is the dipole coupling, E_ν is the SM neutrino energy, $m_{\mathcal{N}}$ is the mass of the heavy neutrino, M is the target mass, $t = -(p_{\mathcal{N}} - p_\nu)^2$ is the momentum transfer, $E_r = -t/2M$ is the target recoil energy, and $F_{1/2}(t)$ are the electric/magnetic target form factors, respectively. Note that the term proportional to $E_r m_{\mathcal{N}}^4$ in the F_1 line only exists for spinless nuclei, and must be replaced for nonzero spin nuclei [127]. In the case of coherent scattering off of a nucleus, F_1 receives a Z^2 enhancement and is therefore dominant over F_2 . Thus, we explicitly neglect the F_2 term for this analysis. In Ref. [27], F_1 was calculated using a dipole parameterization. However, since the dipole approximation tends to underestimate the cross section at high squared momentum transfer Q^2 , Ref. [31] updated to the Fourier-Bessel parameterization,

$$F^{\text{FB}}(Q) = Z^2 N \frac{\sin(QR)}{QR} \sum_n \frac{(-1)^n a_n}{n^2 \pi^2 - Q^2}, \quad (6.6)$$

¹github.com/Harvard-Neutrino/LeptonInjector

where Z is the proton number of the struck nucleus, N is a normalization factor ensuring $F^{\text{FB}}(0) = 1$, and a_n are coefficients obtained from nuclear scattering data [256–259]. For the case of inelastic upscattering off of individual nuclei, we use the nucleon electric and magnetic form factors discussed in the appendix of Ref. [27]. Inelastic upscattering exhibits only a linear enhancement in the number of protons and neutrons, which suppresses its effect compared to the Z^2 enhancement of the coherent upscattering cross section. Ref. [31] also includes contributions from the sub-dominant helicity-conserving channel of both the elastic and inelastic upscattering cross sections.

The total upscattering cross section was obtained by integrating equation (6.5) over the physical momentum transfer range [130]. The overall neutrino production rate via upscattering in the BNB was then calculated using this cross section in combination with the density of each nuclear and nucleon target available along the BNB. These include nuclei in the bedrock and air between the BNB target and the detector, as well as carbon and hydrogen nuclei in the detector itself. The bedrock composition is defined in table 6.1 and is considered to have a density of 2.9 g/cm³. Upscattering in the bedrock (detector) is more important for long-lived (short-lived) neutrinos. The kinematics of the outgoing neutrino were fixed by sampling a momentum transfer t according to equation (6.5), which consequently determines the nuclear recoil energy $E_r = -t/2M$ as well as the neutrino energy and scattering angle [127],

$$\begin{aligned} E_{\mathcal{N}} &= E_{\nu} - E_r, \\ \cos(\theta) &= \frac{E_{\nu} - E_r - ME_r/E_{\nu} - m_{\mathcal{N}}^2/2E_{\nu}}{\sqrt{E_{\nu}^2 + E_r^2 - 2E_{\nu}E_r - m_{\mathcal{N}}^2}}. \end{aligned} \tag{6.7}$$

One can rewrite equation (6.5) to isolate an overall $1/t^2$ prefactor, indicating a preference for $t \approx 0$ and thus $\cos \theta \approx 1$ for this process.

After neutrinos are produced, they will decay via the process $\mathcal{N} \rightarrow \nu\gamma$. The decay width Γ and lab-frame decay length L_{decay} of this process are given by [27,31,

$$\begin{aligned}\Gamma &= \frac{d^2 m_{\mathcal{N}}^3}{4\pi}, \\ L_{\text{decay}} &= 4\pi \frac{\beta E_{\mathcal{N}}}{d^2 m_{\mathcal{N}}^4},\end{aligned}\tag{6.8}$$

where

$$d \equiv \left(\sum_{\alpha} |d_{\alpha N}|^2 \right).\tag{6.9}$$

Considering the dipole parameters preferred by the MiniBooNE excess, the typical neutrino decay length at BNB energies is $L \lesssim 1$ m. The probability of decaying within the MiniBooNE detector is then

$$P_{\text{decay}} = \exp\left(\frac{-L_{\text{enter}}}{L_{\text{decay}}}\right) - \exp\left(\frac{-L_{\text{exit}}}{L_{\text{decay}}}\right),\tag{6.10}$$

where L_{enter} and L_{exit} are the distance along the neutrino path from the upscattering location to the entrance and exit of the MiniBooNE detector, respectively. Assume the neutrinos are Dirac particles, their differential decay which is given by [131, 135]

$$\frac{d\Gamma}{d\cos\theta} = \frac{1}{2}(1 + \alpha \cos\theta),\tag{6.11}$$

where $\alpha = 1$ ($\alpha = -1$) for the decay of a right-handed (left-handed) heavy neutrino, and vice versa for the antineutrino case.

In Ref. [27] and Ref. [31], we used the combination of the production rate from equation (6.5) and the decay rate from equation (6.8) to compute the single photon rate from neutrino decay in MiniBooNE.

6.3.1 Simulation in LeptonInjector

As stated above, Ref. [27] built a custom simulation to assess the neutrino model in MiniBooNE. Ref. [31] instead performed the MiniBooNE simulation using the `LeptonInjector` software package. `LeptonInjector` was originally developed within the IceCube collaboration [260] to simulate the interactions of high-energy neutrinos

Nucleus	Z	A	Crust Mass Fraction	Crust Atomic Fraction	Nuclear Radius [MeV ⁻¹]
O	8	16	0.466	0.627	0.00218
Si	14	28	0.277	0.213	0.00252
Al	13	27	0.081	0.065	0.00247
Fe	26	56	0.05	0.019	0.00301
Ca	20	40	0.037	0.02	0.00281
K	19	39	0.027	0.015	0.00277
Na	11	23	0.026	0.024	0.00241
Mg	12	24	0.015	0.013	0.00247
Ti	22	48	0.004	0.002	0.0029
P	15	31	0.001	0.001	0.00257

Table 6.1: Relevant parameters of the ten most abundant nuclei in the Earth’s upper crust according to [32]

within their detector; however, this version was not suitable to handle the neutrissimo model out-of-the-box. A series of updates were made to `LeptonInjector`, led by Austin Schneider and myself, in order to carry out the study described here. These include (but are not limited to):

- Support for the input of total and differential cross section tables.
- Support for a wide variety of detector subsystem geometries, including extruded polygons.
- Ability to simulate a tree of secondary interactions following the initial interaction.
- Ability to calculate event weights for general interaction trees.
- Support for a combination of cross sections and decays as possible interactions for a given particle.
- Support for a fiducialization requirement to improve the computational efficiency.

The current version of `LeptonInjector` is suitable for general purpose studies of particle interactions in a wide variety of detectors. It was presented at the March 2023 Coherent-CAPTAIN-Mills Workshop [261] and is publicly available on GitHub [262].

While `LeptonInjector` is written in C++, the classes and methods are importable as a library in Python [263].

We now briefly review the `LeptonInjector`-based MiniBooNE analysis presented in Ref. [31]. Neutrinos are first injected along the beam axis with energies sampled from the BNB flux [14]. The upscattering location is then sampled according to the material-dependent interaction length along the neutrino path. The interaction length in each material is calculated via upscattering cross section tables computed using `DarkNews` [264]. In order to optimize the efficiency of the simulation, neutrinos are required to undergo a Priamkoff upscattering interaction within three neutrissimo decay lengths of the MiniBooNE detector. The kinematics of the neutrissimo are fixed by sampling the differential cross section in equation (6.5), calculated in tabular format using `DarkNews`. The decay location of the neutrissimo is then sampled according to the lab frame decay length in equation (6.8). If the neutrissimo path crosses the detector volume, it is required to decay within or just before the detector, further optimizing the efficiency of the simulation. Once a photon is generated, we sample the pair-production location according to the radiation length of each traversed material, which we use to perform a robust fiducialization cut. The weight of each event is calculated by comparing the generation probability of each sampled quantity (e.g. the energy/direction of the initial neutrino, the kinematics of the produced neutrissimo/photon, and the upscattering/decay locations) to the specified physical probability distribution. The flux units provided in Ref. [14] allow us to compute event weights in units of POT^{-1} .

Figure 6-5 shows a schematic depiction of the `LeptonInjector` neutrissimo simulation in MiniBooNE, which incorporates 541 m of bedrock, a 9 m-radius spherical air-filled detector hall, and the 6.1 m-radius spherical mineral oil detector. The result of this procedure is a robust Monte Carlo sample of single photons in MiniBooNE from neutrissimo decay. The kinematic variables of each observable photon have been carefully computed, and the detailed treatment of the geometry of the MiniBooNE detector hall allows for a realistic estimation of the fiducialization requirement. While the latter point is not necessarily vital for the relatively simple MiniBooNE detector,

it is much more important when considering the complex MINER ν A detector for the analysis discussed in section 6.4. The careful treatment of new physics models enabled by `LeptonInjector` is essential for differentiating between the many proposed explanations of MiniBooNE, both when determining the preferred region in model parameter space from MiniBooNE data and when determining constraints from other experiments. This is one of the intended use cases motivating our recent public release of `LeptonInjector` [262].




Figure 6-5: Schematic depiction of the neutrino model in MiniBooNE as simulated using `LeptonInjector`. Figure 6-5a shows the simulation of Primkaoff upscattering along the beamline, and figure 6-5b shows an example of upscattering, neutrino decay and pair-production within the MiniBooNE detector.

6.3.2 Fits to the MiniBooNE Excess

After performing the procedure outlined in section 6.3.1, we down-weighted each event according to a linear parameterization of the MiniBooNE photon detection efficiency [27, 265]. We also imposed a reconstruction threshold on the true photon kinetic energy requiring $E_\gamma > 140$ MeV. The energy (angle) of the photon was smeared according to a power-law (quadratic) fit to a MiniBooNE single electromagnetic shower simulation sample [266]. The typical resolution in the energy (angular)

distribution was 10% (3°) for the photons generated in this model, consistent with values reported by MiniBooNE [267].

After these steps, the simulated events could then be used to estimate the single photon rate from neutrino decays in the E_ν^{QE} and $\cos\theta$ distributions of MiniBooNE’s electron-like sample. We performed a separate fit to the MiniBooNE excess in each distribution. After subtracting the SM prediction and the oscillation prediction from the global fit without MiniBooNE, we calculated a χ^2 test statistic between the data and the neutrino prediction throughout a grid of points in $d_{\mu\mathcal{N}}-m_{\mathcal{N}}$ parameter space. For the E_ν^{QE} distribution, we used the MiniBooNE electron-like covariance matrix provided in the data release [268] associated with Ref. [10]. The MiniBooNE collaboration has not released any information regarding systematic uncertainty in the $\cos\theta$ distribution; therefore, we assign an uncorrelated fractional systematic uncertainty of 13% in this distribution, consistent with the overall level in the E_ν^{QE} distribution. We compute confidence regions in neutrino parameter space using the $\Delta\chi^2$ test statistic, assuming Wilks’ theorem [26] for two degrees of freedom.

Figure 6-6 shows the 95% and 3σ confidence level allowed regions in neutrino parameter space from the MiniBooNE E_ν^{QE} and $\cos\theta$ distributions. These allowed regions come from Ref. [31], which is intended to supersede the results in Ref. [27]. This is because Ref. [31] makes a number of improvements upon the analysis in Ref. [27], including the use of a more robust nuclear electromagnetic form factor and a more detailed simulation within `LeptonInjector`. One can see from figure 6-6 that the energy and angular distributions of the MiniBooNE excess prefer different regions of parameter space. This is because the angular distribution has a not-insignificant component at large angles, $-1 \leq \cos\theta \leq 0$, as shown in figure 2-6b. Larger neutrino masses are required to explain this region, as such neutrinos will have a lower Lorentz boost and thus result in a photon with broader lab-frame angles with respect to the beamline. This is in contrast with the energy distribution, which prefers lower neutrino masses such that the lower energy part of the BNB neutrino flux is not kinematically forbidding from upscattering into a neutrino. The difference between preferred regions of the energy and angular distributions in such visible heavy neu-

trino decay models has been discussed previously in the literature [128,135]; however, our study is the first comprehensive analysis of the compatibility of both distributions. Our results in figure 6-6 indicate an overlap between the allowed regions of the E_ν^{QE} and $\cos\theta$ allowed regions at the 95% C.L. Figure 6-7 shows comparisons between the MiniBooNE excess in each distribution and the mixed model prediction, considering the neutrissimo model indicated by the black star in figure 6-6, which is near this 95% C.L. overlap. One can see that the oscillation contribution to the MiniBooNE LEE in this model is indeed small compared with the neutrissimo contribution. A larger oscillation contribution might help explain the $-1 \leq \cos\theta \leq 0$ region of the excess—this point is discussed further in section 6.4. That being said, the mixed model does a better job of predicting the excess in the lowest-energy region. This can be seen by comparing the compatibility between data and prediction in the first two bins of figure 6-7a to that in figure 2-5a.

There are a number of existing constraints on the neutrissimo model in the $10 \leq m_{\mathcal{N}} [\text{MeV}] \leq 1000$ regime, as indicated by the grey regions in figure 6-6. These include Super-Kamiokande [233,237], Borexino [233], Supernova 1987A [81,130], LSND [130], CHARM-II [232], and NOMAD [231]. See section 6.4 for a more detailed discussion on each constraint. We briefly note that the overlap in the 95% allowed regions in MiniBooNE falls close to the NOMAD constraint derived in Ref. [231], which may motivate a re-analysis of the NOMAD data.

Finally, we remind the reader that any explanation of MiniBooNE must be compatible not only with the energy and angular distributions of the excess but also with the timing distribution [10]. Specifically, the new physics events cannot experience time delays $\Delta t \gtrsim 10$ ns. This was studied within the context of the neutrissimo model in Ref. [27]. As shown in figure 6-8, neutrissimos with $m_{\mathcal{N}} \sim 500$ MeV have time delays $\Delta t \sim 1$ ns, certainly within the constraint imposed by the MiniBooNE timing distribution. This is because in the region of parameter space preferred by MiniBooNE, neutrissimos tend to be very short-lived and must therefore be produced via upscattering off of nuclei in the detector. As such, the total travel time is dominated by the original speed-of-light neutrino.




Figure 6-6: Allowed regions at the 95% and 3σ confidence level in $d_{\mu\mathcal{N}}-m_{\mathcal{N}}$ obtained through fits to the MiniBooNE excess in the E_{ν}^{QE} and $\cos\theta$ distributions. Existing 2σ constraints on this model are indicated by the grey regions.

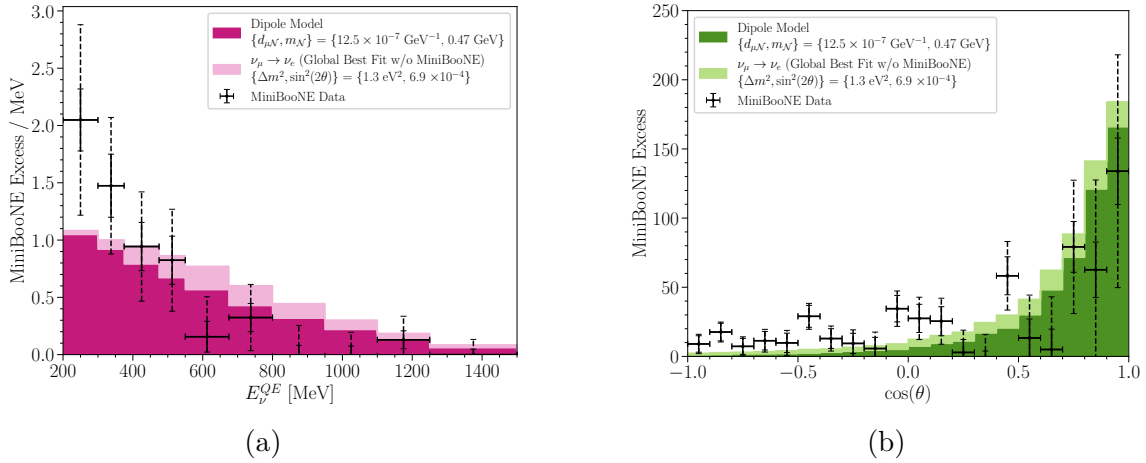


Figure 6-7: Figure 6-7a and figure 6-7b show the E_{ν}^{QE} and $\cos\theta$ distributions of the MiniBooNE excess, respectively, compared with the prediction from the neutrino model indicated by the black star in figure 6-6. The oscillation contribution from the $3 + 1$ global fit without MiniBooNE is also shown.




Figure 6-8: The added time delay in MiniBooNE for a neutrissimo with the indicated parameters, as calculated in Ref. [27].

6.4 Publication: *Dipole-coupled neutrissimo explanations of the MiniBooNE excess including constraints from MINERνA data*

We close this chapter with a presentation of Ref. [31]. The work was led by myself, Matheus Hostert, and Austin Schneider. The full *Physical Review D* publication is included below. As described in section 6.3, this study performed a more detailed analysis of the mixed model described in section 6.2 as a solution to the MiniBooNE excess. We also derived world-leading constraints on the neutrissimo model using neutrino electron elastic scattering data from the MINERνA experiment [217–219]. While these constraints are strong, they do not rule out the overlap in the 95% C.L. allowed regions from the MiniBooNE energy and angular distributions shown in figure 6-6. As shown in Figure 8 of the paper below, this is a result of the strict kinematic cuts employed in the MINERνA elastic scattering analysis. If these cuts are relaxed, MINERνA should be sensitive to the MiniBooNE neutrissimo solution in the mixed model. The most important result from this study is given in Figure 10, which shows the surviving parts of the MiniBooNE allowed regions in neutrissimo parameter space after the MINERνA constraints are incorporated. The dashed lines in the right

panel indicate contours of constant number of neutrino decays in the MINER ν A fiducial volume before any cuts are applied. One can see that MINER ν A should have $10^4 - 10^5$ events before cuts within the MiniBooNE-preferred region of parameter space. We also test the implications of weak-flavor-conserving UV-completions of the dipole operator of equation (6.1), which predict $d_{\alpha N} \propto m_{\ell_\alpha}$ and thus a large d_τ . This scenario mostly alters the allowed regions and constraints in the lower mass region of parameter space without impacting the MiniBooNE solution.

Note: while the published version of this article uses the phrase “heavy-neutral-lepton” in the title, the authors prefer the arXiv title given above.

Dipole-coupled heavy-neutral-lepton explanations of the MiniBooNE excess including constraints from MINERvA data

N. W. Kamp¹, M. Hostert^{2,3}, A. Schneider¹, S. Vergani⁴,
C. A. Argüelles⁵, J. M. Conrad¹, M. H. Shaevitz⁶, and M. A. Uchida⁴

¹*Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

²*University of Minnesota, Minneapolis, Minnesota 55455, USA*

³*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2J 2W9, Canada*

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

⁵*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

⁶*Department of Physics, Columbia University, New York, New York 10027, USA*



(Received 9 July 2022; accepted 15 February 2023; published 9 March 2023)

We revisit models of heavy neutral leptons (neutrissimos) with transition magnetic moments as explanations of the 4.8σ excess of electronlike events at MiniBooNE. We first reexamine the preferred regions in the model parameter space to explain MiniBooNE, considering also potential contributions from oscillations due to an eV-scale sterile neutrino. We then derive constraints on the model using neutrino-electron elastic scattering data from MINERvA. To carry out these analyses, we have developed a detailed Monte Carlo simulation of neutrissimo interactions within the MiniBooNE and MINERvA detectors using the `LeptonInjector` framework. This simulation allows for a significantly more robust evaluation of the neutrissimo model compared to previous studies in the literature—a necessary step in order to begin making definitive statements about beyond the Standard Model explanations of the MiniBooNE excess. We find that MINERvA rules out a large region of parameter space, but allowed solutions exist at the 2σ confidence level. A dedicated MINERvA analysis would likely be able to probe the entire region of preference of MiniBooNE in this model.

DOI: [10.1103/PhysRevD.107.055009](https://doi.org/10.1103/PhysRevD.107.055009)

I. INTRODUCTION

Past, current, and future neutrino experiments offer some of the most promising avenues for observing physics beyond the Standard Model (BSM). The discovery of neutrino oscillations, and thus nonzero neutrino mass, is itself an indication of BSM physics [1,2], spurring a decades-long oscillation experimental program spanning orders of magnitude in energy and length scales [3–5]. A nearly-consistent three-neutrino mixing paradigm has emerged from this program; however, anomalous results have also been observed. Two striking examples are the excess of inverse-beta-decay events at the LSND detector at the Los Alamos Neutrino Science Center [6] and the excess of electronlike events observed by the MiniBooNE (MB) experiment at the Fermilab Booster Neutrino Beam (BNB) [7]. Determining the nature of these excesses is an active frontier in neutrino physics [8]. Historically, both

excesses have been interpreted within the context of a $3 + 1$ model, in which one introduces an eV-scale sterile neutrino facilitating short-baseline $\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_e(\bar{\nu}_e)$ oscillations. However, models beyond the vanilla $3 + 1$ scenario may better accommodate these anomalies within the global experimental landscape, including cosmology and other neutrino oscillation experiments.

One such model considers MeV-scale heavy neutral leptons (HNLs), or “neutrissimos” [9,10], with transition magnetic moments, popularly known as the dipole portal to HNLs [11–24]. Most recently, Ref. [23] showed that the MB anomaly could be fitted with such a dipole portal extended to include a mass-mixed eV-scale sterile neutrino. The advantages of this model are twofold: (1) it provides a better description of the low-energy, forward-angle part of the MB excess compared to the $3 + 1$ model alone, and (2) reducing the oscillation-based contribution to the MB excess alleviates tension in global fits to the $3 + 1$ model while retaining an explanation to the LSND anomaly.

For the main result of this paper, we derive constraints on the dipole portal coupling using existing experimental results from the MINERvA Collaboration. Specifically, MINERvA has performed measurements of the neutrino-electron elastic scattering ($\nu - e$) rate for the purpose of

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. 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.

constraining the Neutrino Main Injector (NuMI) low-energy (LE) and medium-energy (ME) neutrino fluxes [25–27]. Photons produced in the decay of the dipole-coupled neutrino would mimic the single electromagnetic shower morphology of electron-scattering (ES) events and would therefore enter as a photonlike (large dE/dx) background in the MINERvA analysis. In addition to deriving new constraints from MINERvA, we improve upon the MiniBooNE dipole model analysis performed in Ref. [23].

The analyses described in this article have been carried out using a novel simulation developed within the `LeptonInjector` framework [28]. This tool allows a more robust description of neutrino interactions in detector subsystems as well as the position and kinematics of observable final state particles. This is an important improvement over previous treatments in the literature, as an accurate simulation of exotic BSM physics scenarios in different neutrino detectors will be vital in determining the nature of the MiniBooNE excess.

The rest of this article is organized as follows. In Sec. II, we review the dipole-portal sterile neutrino model in more detail. In Sec. III, we introduce the novel `LeptonInjector`-based simulation developed for the studies presented here. In Sec. IV, we refine the preferred regions in dipole parameter space which explain the energy and angular distributions of the MiniBooNE excess. In Sec. V, we calculate constraints in dipole parameter space derived from the MINERvA ES analysis. In Sec. VI, we discuss existing and projected constraints on the dipole model from current and future neutrino experiments. Finally, in Sec. VII we discuss the outlook of this mixed model of oscillation and decay as a solution to the MiniBooNE anomaly in light of our derived MINERvA constraints.

II. THE DIPOLE PORTAL

We introduce a right-handed neutrino \mathcal{N} that couples to the left-handed neutrino fields via a transition magnetic moment. At the effective operator level, we have the dimension-six dipole operators,

$$\mathcal{L} \supset \frac{1}{\Lambda^2} \bar{L}_\alpha \tilde{H} \sigma^{\mu\nu} \mathcal{N}_R (C_B^\alpha B_{\mu\nu} + C_W^\alpha W_{\mu\nu}^a \sigma_a) + \text{H.c.}, \quad (1)$$

where α is a flavor index, $\tilde{H} = i\sigma_2 H^*$, C_B^α , and C_W^α are Wilson coefficients and Λ is the new physics scale. After electroweak symmetry breaking, the dipole operator gives rise to the electromagnetic transition magnetic moment of neutrinos,

$$\mathcal{L} \supset d_{\alpha\mathcal{N}} \bar{\nu}_\alpha \sigma_{\mu\nu} F^{\mu\nu} \mathcal{N}_R + \text{H.c.}, \quad (2)$$

where ν_α corresponds to the neutrino Weak eigenstates and $F^{\mu\nu}$ to the electromagnetic field strength. The dipole parameter is defined as $d_{\alpha\mathcal{N}} = (v_h/\sqrt{2})(C_W^\alpha C_B^\alpha + C_{W,S}^\alpha)/\Lambda^2$, where v_h is the vacuum expectation value of

the SM Higgs. The other transition moments mediated by the W and Z bosons will also be present, but their low-energy effects are further suppressed by G_F , and therefore negligible in our region of interest.

The upscattering signature we are interested in at MiniBooNE is initiated by muon neutrinos and antineutrinos, so in this work, we always consider the second-generation coupling $d_{\mu\mathcal{N}}$. In UV completions of this operator from one-loop diagrams, the size of the dipole coupling is typically proportional to $d_{\alpha\mathcal{N}} \sim (1/16\pi^2)(m_\beta/m_X^2)$, where m_β is the mass of some charged particle and X is some heavy, charged scalar, for example. Under the assumption of flavor-conserving interactions between SM neutrinos and the new physics, one would take $\beta = \alpha$ and conclude that the transition magnetic moment of the third-generation neutrinos is much larger. While this need not necessarily be the case, it still provides enough motivation for us to consider the third-generation coupling $d_{\tau\mathcal{N}}$, including the case

$$\frac{d_{\tau\mathcal{N}}}{d_{\mu\mathcal{N}}} = \frac{m_\tau}{m_\mu}, \quad (3)$$

and neglecting the first-generation couplings altogether.

We note that large neutrino magnetic moments typically imply large Dirac masses for light neutrinos, $m_D \bar{\nu}_L \mathcal{N}_R$ [29–31]. In models with heavy neutrinos, this presents two challenges: (i) neutrino masses, schematically given by $m_\nu \sim m_D^2/M$, may be too large, and (ii) m_D will generate large mixing between active and heavy neutrinos, $U_{\alpha\mathcal{N}} \sim m_D/M$. Here, M stands for the heavy neutrino mass scale. The first challenge is easily overcome in models like the inverse-seesaw where lepton number is approximately conserved and right-handed neutrino fields combine into pseudo-Dirac \mathcal{N} particles. In that case, m_ν is controlled by the mass splitting between the Majorana neutrinos, which may be parametrically small. This is also the case preferred by the short-baseline phenomenology discussed below, since the single-photons produced in the decays of Dirac HNLs are less forward. The second challenge, however, is not so easy to overcome. The mixing between active and heavy neutrinos remains large even in inverse-seesaw models and the parameter space in that case is strongly constrained by laboratory limits on $U_{\alpha\mathcal{N}}$. For instance, decay-in-flight signatures at neutrino experiments, where \mathcal{N} is copiously produced in meson decays at the target, can set limits as strong as $|U_{\alpha\mathcal{N}}|^2 < \mathcal{O}(10^{-11})$ [32].

There are several models in the literature that can suppress the Dirac mass in comparison with the magnetic moment of neutrinos [33–39], but only to a certain extent and not without fine tuning. We proceed assuming that m_D is sufficiently small so as not to impact the phenomenology, but note that depending on the amount of fine tuning, bounds on the mixing angles $|U_{\alpha\mathcal{N}}|$ would also need to be considered.

In performing a fit to MiniBooNE data, we also include a sterile neutrino, ν_s , which mixes with light neutrinos but does not have transition magnetic moments with active neutrinos. This sterile will be responsible for short-baseline oscillations with Δm_{41}^2 of $\mathcal{O}(1 \text{ eV}^2)$. We provide two examples, the case of a global fit to short-baseline data excluding MiniBooNE [23] and the case of a joint fit to only MiniBooNE and the recent MicroBooNE CCQE-like analysis [40]. It is also important that ν_s and the heaviest neutrino do not mix, as otherwise \mathcal{N}_R would also mix with light neutrinos via ν_s . In summary, our spectrum is defined as

$$\nu_i \sim \sum_{\alpha=\{e,\mu,\tau,s\}} U_{\alpha i} \nu_\alpha, \quad \text{for } i \leq 4 \quad \text{and} \quad \nu_5 \equiv \mathcal{N}, \quad (4)$$

where the approximate symbol means that any additional term contains very small mixing elements. In addition, this mixing would result in a corresponding dipole coupling $d_{\alpha s} \bar{\nu}_\alpha \sigma_{\mu\nu} F^{\mu\nu} \nu_s$, which is strongly constrained by big bang nucleosynthesis and stellar cooling. While the former may be modified *à la* secret interactions, the latter is significantly more robust and constrains $d_{\mu s} < 2 \times 10^{-12} \mu_B$ [41].

The relevant interactions for this work are shown in Fig. 1, which include Primakoff upscattering off of a nuclear target $\nu A \rightarrow \mathcal{N} A$ (left) and the radiative decay of the heavier neutrino $\mathcal{N} \rightarrow \nu_\alpha \gamma$ (right). Due to the photon propagator, the scattering process is dominated by coherent exchange with the nucleus, except at the largest $m_{\mathcal{N}}$ values. We include helicity-flipping upscattering, where the helicity of \mathcal{N} is opposite to that of ν_α , and helicity-conserving, where they are the same. The latter is suppressed by the typical energy of the process, $m_{\mathcal{N}}^2/E^2$, and is only relevant at the largest masses. To give an example, the cross section for an incoming neutrino with energy $E_{\nu_\alpha} = 1 \text{ GeV}$ upscattering off of a carbon nucleus into a neutriissimo with mass $m_{\mathcal{N}} = 50 \text{ MeV}$ and dipole coupling $d_{\alpha \mathcal{N}} = 10^{-6} \text{ GeV}^{-1}$ is $\sigma \approx 5 \times 10^{-39} \text{ cm}^2$ [18].

The rest frame decay width of the \mathcal{N} is given by the incoherent sum over all outgoing flavors,

$$\Gamma_{\mathcal{N} \rightarrow \nu \gamma} = \left(\sum_{\alpha} |d_{\alpha \mathcal{N}}|^2 \right) \frac{m_{\mathcal{N}}^3}{4\pi}. \quad (5)$$

Throughout this work, we will assume that either \mathcal{N} or light neutrinos are Dirac particles, such that the differential decay




FIG. 1. \mathcal{N} production from ν upscattering (left) and \mathcal{N} decay (right).

rate is proportional to $(1 \pm \cos \theta)$ [19]. For the region of dipole parameter space preferred by MiniBooNE [23], the HNL has a lab frame decay length $L \sim 1\text{--}10 \text{ m}$ for typical MINERvA neutrino energies and $L \leq 1 \text{ m}$ for typical MiniBooNE neutrino energies. Thus, Primakoff upscattering near or within each detector is the most relevant production mechanism in this region of parameter space. For smaller HNL masses, decay lengths become longer and upscattering in the dirt along the BNB and NuMI beamlines becomes more important for MiniBooNE and MINERvA, respectively.

III. SIMULATING NEUTRISSIMOS

In order to describe neutriissimo interactions in the MiniBooNE and MINERvA detectors, we have developed a custom simulation based in the `LeptonInjector` framework [28]. The simulation begins by injecting neutrinos in MiniBooNE and MINERvA according the BNB and NuMI beam profiles, respectively. For each neutrino, a flight path is randomly selected within a cone surrounding the detector. This path is then used to calculate intersections with different components of the beamline, including bedrock between the target and detector as well as various detector subsystems. An upscattering location is sampled along this flight path according to the cross section in each traversed material. The final state kinematics of the produced HNL are sampled according to the $d\sigma/dy$ distribution, where $y \equiv (E_\nu - E_{\mathcal{N}})/E_\nu$ in the lab frame [41]. See the Appendix for more details on the upscattering cross section used in this analysis. Next, a decay location is sampled along the flight path of the HNL. The final state kinematics of decay photon are sampled according to $d\Gamma/d\cos\theta \propto (1 - \cos\theta)$. We trace the flight path of the photon through different materials until it converts to an e^+e^- pair, which are assumed to reconstruct as a single electromagnetic shower within the MiniBooNE and MINERvA detectors.

The above procedure is easily generalizable to other BSM scenarios and detector configurations. The user needs only to provide (1) the flux of initial state particles (neutrinos in this case), (2) the total and differential cross sections for the relevant processes, and (3) the relevant detector geometry. The flux and cross section can be input either as analytic expressions or splines; `LeptonInjector` will interpolate in the latter case. The detector geometry is set using a simple configuration file in which the user instantiates any number of volumes. `LeptonInjector` supports most simple volumes as well as general extruded polygons—for example, the nuclear targets in MINERvA shown in Fig. 2, which are subsections of hexagonal prisms. `LeptonInjector` will then generate user-specified initial state particles, sample their interaction locations, and store the kinematics and weights of the final state particles. The simulation is set up to be as efficient as possible, such that most generated events create an observable final state within

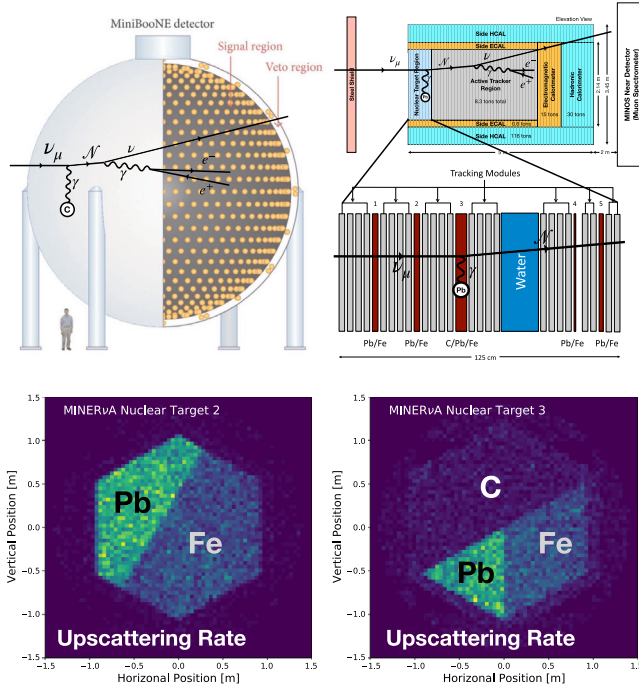


FIG. 2. (Top) Schematic representation of HNL production via upscattering and the subsequent HNL decay within the MiniBooNE (left) and MINERvA (right) detectors. Detector images have been adapted from Refs. [46,47]. (Bottom) Example upscattering rates within two of the MINERvA nuclear targets as simulated using *LeptonInjector*. The coherent enhancement of the upscattering cross sections leads to a larger rate in the high- Z components.

the fiducial volume and are appropriately downweighted. In the analysis presented here, we have tracked final state particles until the production of an observable e^+e^- pair. This strategy can be adapted for other BSM scenarios which involve multiple interactions before the production of an observable final state.

The simulation performed for the analysis presented here gives a robust estimation of the e^+e^- event rate within the fiducial volume of each experiment. Further, it provides an accurate kinematic description of the photons which survive fiducialization (namely, the photon energy and angle with respect to the beamline). This is vital, as it will be shown that the region in dipole parameter space preferred by the MiniBooNE excess is highly sensitive to the kinematics of the final state photon. Such an effect has been appreciated by previous studies in the literature [11,24], but a comprehensive fit over the full dipole parameter space was not attempted until Ref. [23]; the simulation described above allows us to refine the fit from Ref. [23].

The simulation is also vital in order to properly assess the ability of MINERvA to constrain the dipole solution to the MiniBooNE excess. By leveraging our simulation, we will show that the reconstruction efficiency for radiative HNL decays in MINERvA varies by 3–4 orders of magnitude

across the parameter space. Previous calculations in the literature have assumed a constant 10% reconstruction efficiency for radiative HNL decays in MINERvA [42], and have thus incorrectly concluded that MINERvA data rules out the dipole-portal MiniBooNE solution. Further, they do not simulate the complex subcomponents of the MINERvA detector—this is an important step, as upscattering in the high- Z nuclear targets just upstream of the fiducial volume can contribute significantly to the single shower sample for short-lived HNLs.

Thus, the *LeptonInjector*-based simulation performed for this study is essential to evaluate the status of the dipole-coupled neutrissimos as an explanation for the MiniBooNE excess. As the community turns toward more exotic BSM explanations of the MiniBooNE excess, it is imperative that these models are evaluated within the context of realistic detector descriptions. Our simulation framework is the perfect tool for such a task; though the version used in this article was developed specifically to study neutrissimo interactions in MiniBooNE and MINERvA, it can be adapted easily to accommodate additional BSM scenarios and neutrino detectors.

IV. NEUTRISSIMOS AT MiniBooNE

The MiniBooNE detector uses Cherenkov light to detect final state particles produced in neutrino interactions. As electrons and photons both produce electromagnetic showers which show up as distorted Cherenkov rings in the detector, the two particles are indistinguishable in MiniBooNE. Thus, photons from the dipole model could contribute to the MiniBooNE electronlike excess. In the case of a nonzero effective dipole coupling, $d_{\mu N}$, muon neutrinos from Fermilab’s Booster Neutrino Beam will undergo Primakoff upscattering into HNL states. This can happen via photon exchange with nuclear targets both within the dirt between the BNB target and the detector and within the CH_2 detector volume itself.

A. Simulation

As described in Sec. III, we use *LeptonInjector* to simulate the production and decay of HNLs in MiniBooNE [28]. A schematic depiction of this process is shown in the top panel of Fig. 2. Muon neutrinos are injected according to the BNB flux and allowed to upscatter to HNLs along the 541 m baseline between the BNB target and the detector or within the 818-ton CH_2 detector itself. We describe the MiniBooNE detector as a sphere of CH_2 with a total radius of 6.1 m and a fiducial radius of 5 m. The MiniBooNE detector sits within a sphere of air with a radius of 9 m, meant to represent the detector hall. To simulate the BNB, muon neutrinos generated in *LeptonInjector* propagate through 541 m of dirt before reaching the MiniBooNE detector. The neutrino upscattering and subsequent HNL decay are

simulated according to the procedure outlined in Sec. III. We parametrize the postfiducialization photon detection efficiency as a linearly decreasing function of the true photon kinetic energy [23]. We also impose a reconstruction threshold on the photon kinetic energy of 140 MeV [43]. We also independently smear the reconstructed visible energy and scattering angle of each photon according to the resolution of each as a function of true photon kinetic energy. The energy (angular) resolution comes from a power-law (quadratic) fit to simulated single electromagnetic shower events in MiniBooNE [44], with typical values of 10% (3°) for photons from this model, consistent with figures reported by the MiniBooNE Collaboration [45].

B. Analysis methodology

The dipole model has previously been studied as a potential solution to the MiniBooNE anomaly [11–18,22–24]. Specifically, we expand upon the study performed in Ref. [23], which examined a mixed model consisting of an eV-scale ν_4 facilitating short-baseline $\nu_\mu \rightarrow \nu_e$ oscillations as well as an MeV-scale HNL decaying to a photon via the dipole portal mechanism. Ref. [23] found a preference for a dipole-coupled HNL with $m_{\mathcal{N}} \sim 400$ MeV and $d_{\mu\mathcal{N}} \sim 3 \times 10^{-7}$ GeV $^{-1}$.

We make a number of improvements to this analysis. First, we make use of a more robust nuclear electromagnetic form factor, implementing a data-driven Fourier-Bessel function parametrization [48–50] with the data files made available in Ref. [51]. This should be compared with the simpler dipole parametrization, which overestimates the differential cross sections at larger momentum exchange (see the Appendix for more details). This has an impact on the allowed regions at larger heavy neutrino masses, as the form factor used in this study drops off much more quickly at larger Q^2 . It also reduces the contribution from the dipole model at large scattering angles, thereby making it difficult to explain the backscattered lepton angular distribution of the MiniBooNE excess. This effect has been pointed out in previous studies of the dipole model in MiniBooNE [16]. In this work, we perform a more detailed analysis of the dipole parameter space to determine whether solutions exist which can accommodate both the energy and angular distributions of the excess. The statistical treatment for each distribution has been improved—we now consider correlated systematic errors in the reconstructed E_ν^{QE} distribution from the provided covariance matrix (after constraining with the covariance matrix for MiniBooNE’s ν_μ dataset). We also introduce an uncorrelated systematic error of 13% in the $\cos\theta$ distribution, consistent with that in the E_ν^{QE} distribution.

We perform fits to the excess only in neutrino-mode data, as MiniBooNE has collected about an order of magnitude more events in this beam configuration compared to their antineutrino-mode data. We use the simulated photon events

from the above procedure to perform two different spectral analyses across dipole parameter space: one in the E_ν^{QE} distribution and one in the $\cos\theta$ distribution. In both cases, we calculate a χ^2 test statistic comparing the dipole model prediction to the remaining excess after subtracting off the oscillation contribution from the MiniBooNE-less global fit reported in Ref. [23]. In the E_ν^{QE} fit, we use the electronlike channel fractional covariance matrix provided by the MiniBooNE Collaboration after constraining with the covariance matrix in the muonlike channel. No systematic errors are provided for the $\cos\theta$ distribution; therefore, as mentioned above, we consider an uncorrelated fractional systematic error of 13% in each bin of the $\cos\theta$ prediction, consistent with the level in the E_ν^{QE} channel. Confidence regions are drawn using a $\Delta\chi^2$ test statistic, assuming Wilks’ theorem with two degrees of freedom [52].

C. Results

The result from the fit procedure described above is shown in Fig. 3. One can see that it is difficult to explain the E_ν^{QE} and $\cos\theta$ distributions through the same dipole-coupled HNL, as the two distributions prefer different regions of dipole parameter space. The preferred regions overlap at the 2σ C.L., though some of this overlap region is in tension with constraints derived from the NOMAD single-photon analysis [53]. As stated above, the difference between this result and the result in Ref. [23] is driven mainly by the updated form factor. This reduces the Primakoff upscattering rate at large scattering angles,




FIG. 3. The 2σ and 3σ C.L. preferred regions to explain the MiniBooNE anomaly in mass-coupling parameter space for a dipole-coupled heavy neutral lepton. The pink (green) curves correspond to results from fitting the E_ν^{QE} ($\cos\theta$) distribution. Dipole model fits are performed after subtracting the oscillation component from a global fit to a $3+1$ model excluding MiniBooNE data [23]. In general, the energy and angular distributions prefer different regions of parameter space, though overlap exists at the $2\sigma - 3\sigma$ level.

requiring the fit to push to larger \mathcal{N} masses and dipole couplings to explain this part of the MiniBooNE excess. That being said, there is a region in parameter space for $d \sim 10^{-6}$ GeV $^{-1}$ and $m_{\mathcal{N}} \sim 0.5$ GeV which is (1) within the 2σ allowed region for the E_{ν}^{QE} distribution, (2) within the 2σ allowed region for the $\cos\theta$ distribution, and (3) not ruled out by NOMAD's single-photon search [53]. In Fig. 4 we show E_{ν}^{QE} and $\cos\theta$ distributions for an example parameter point in this region, located at $d = 1.25 \times 10^{-6}$ GeV $^{-1}$ and $m_{\mathcal{N}} = 0.47$ GeV. One can see that this model can describe most of the excess except for the region at $\cos\theta \lesssim 0$.

This situation might be improved when considering alternative oscillation scenarios. The above fits assumed the MiniBooNE-less global-fit result, which found a best-fit solution at $\{\Delta m^2 \approx 1.3$ eV 2 , $\sin^2(2\theta_{\mu e}) \approx 6.9 \times 10^{-4}\}$ [23]. We now consider an alternative sterile neutrino hypothesis:




FIG. 4. The E_{ν}^{QE} (top) and $\cos\theta$ distributions at the example dipole model hypothesis indicated by the black star in Fig. 3. The darker contribution in each stacked histogram corresponds to the dipole model prediction, while the lighter contribution corresponds to the oscillation contribution. The background-subtracted MiniBooNE excess is indicated by the black data points, with solid and dashed error bars indicating statistical and statistical + systematic errors, respectively.

the result from the recent MiniBooNE + MicroBooNE CCQE-like 3 + 1 combined fit performed by the MiniBooNE Collaboration [40].

The MiniBooNE + MicroBooNE CCQE-like 3 + 1 combined analysis found a best-fit solution at $\{\Delta m^2 \approx 0.2$ eV 2 , $\sin^2(2\theta_{\mu e}) \approx 0.03\}$ [40]. This introduces a much larger $\nu_{\mu} \rightarrow \nu_e$ oscillation component in MiniBooNE. Thus, the dipole model is primarily driven to explain the lowest energy and most forward-angle portion of the excess. A mild preference for a dipole-coupled heavy neutral lepton is found at the 1σ level. As the dipole model is no longer required to explain the broad-angle portion of the MiniBooNE excess, the angular fit is able to accommodate a large range of heavy neutrino masses while the energy fit prefers lower heavy neutrino masses at $m_{\mathcal{N}} \lesssim 100$ MeV. The preferred regions in dipole parameter space under this oscillation hypothesis are shown in Fig. 5. As a benchmark point, we consider a solution at $d = 1.7 \times 10^{-7}$ GeV $^{-1}$ and $m_{\mathcal{N}} = 0.08$ GeV. As shown in Fig. 6, this benchmark point can reasonably describe the E_{ν}^{QE} and $\cos\theta$ distributions of the MiniBooNE excess.

One can also consider a nonzero transition magnetic moment coupling between the \mathcal{N} and the ν_{τ} flavor eigenstate. This would open up the decay channel $\mathcal{N} \rightarrow \nu_{\tau}\gamma$, increasing the decay width by the ratio $(|d_{\mu\mathcal{N}}|^2 + |d_{\tau\mathcal{N}}|^2)/|d_{\mu\mathcal{N}}|^2$. This will have a more pronounced impact on the fit in the lower HNL mass region of parameter space, as lifetimes in the higher HNL mass region are sufficiently




FIG. 5. 1σ C.L. preferred and 2σ C.L. allowed regions with regard to the MiniBooNE anomaly in mass-coupling parameter space for a dipole-coupled heavy neutral lepton. The pink (green) curves correspond to results from fitting the E_{ν}^{QE} ($\cos\theta$) distribution. Dipole model fits are performed after subtracting the oscillation component from the combined MiniBooNE + MicroBooNE CCQE-like 3 + 1 fit [40]. Mild preference for a dipole-coupled heavy neutrino with $m_{\mathcal{N}} \lesssim 100$ MeV is found at the 1σ level.




FIG. 6. Similar to Fig. 4. The E_{ν}^{QE} (top) and $\cos\theta$ (bottom) distributions at the example dipole model hypothesis indicated by the black star in Fig. 5. The oscillation component comes from the combined MiniBooNE + MicroBooNE CCQE-like 3 + 1 fit [40].

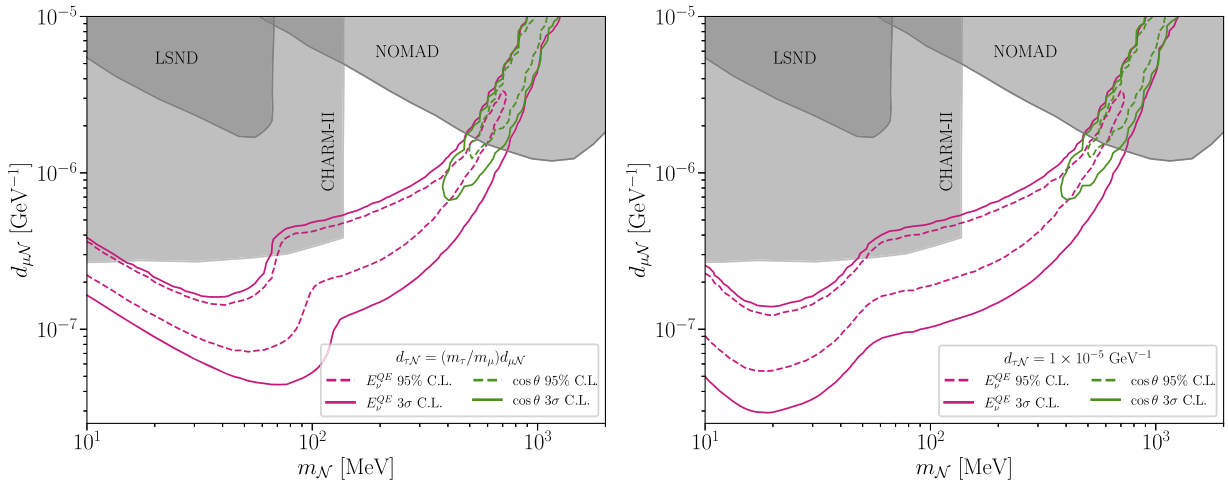


FIG. 7. Similar to Fig. 3, but considering the $\mathcal{N} \rightarrow \nu_{\tau}\gamma$ decay channel. On the left, we show the case with coupling $d_{\tau\mathcal{N}} = (m_{\tau}/m_{\mu})d_{\mu\mathcal{N}}$ and on the right we consider $d_{\tau\mathcal{N}} = 1 \times 10^{-5} \text{ GeV}^{-1}$ instead. Compared to Fig. 3, preferred regions in the lower HNL mass region move to lower $d_{\mu\mathcal{N}}$ values.

short such that introducing another decay channel does not appreciably change the phenomenology.

As discussed in Sec. II, in some UV completions of the dipole model, a natural scaling given by $d_{\tau\mathcal{N}}/d_{\mu\mathcal{N}} = m_{\tau}/m_{\mu}$. The resulting preferred regions in dipole model parameter space under this assumption are shown on the left panel of Fig. 7. One can see that, compared with Fig. 3, solutions explaining the E_{ν}^{QE} distribution have opened at lower $d_{\mu\mathcal{N}}$ couplings for $m_{\mathcal{N}} \lesssim 100 \text{ MeV}$. We also examine the effect of large tau coupling $d_{\tau\mathcal{N}} = 1 \times 10^{-5} \text{ GeV}^{-1}$, which is meant to capture the extent of flexibility introduced into the dipole model when allowing for nonzero $d_{\tau\mathcal{N}}$. The preferred regions for this case are shown in the right panel of Fig. 7. For both cases, we consider an oscillation contribution given by the MiniBooNE-less global fit.

V. NEUTRINOS AT MINERvA

Neutrino upscattering can also occur in the MINERvA detector. We choose to study MINERvA for two main reasons: (i) the NuMI beam is a higher-energy beam in comparison with the BNB. This is specially true for the medium-energy (ME) NuMI configuration, where $\langle E_{\nu} \rangle \simeq 7 \text{ GeV}$, but it is still the case for the low-energy (LE) configuration, where $\langle E_{\nu} \rangle \simeq 3 \text{ GeV}$. This allows us to probe HNLs of larger masses. (ii) it is one of the few accelerator experiments in the few GeV region to have a dedicated neutrino-electron ($\nu - e$) scattering analysis. While measurements of this channel have been performed with greater precision at experiments like CHARM [54] and CHARM-II [55], LSND [56], reactors [57], Borexino [58,59], and Super-Kamiokande [60], they are not as well suited for the study of the MiniBooNE explanations considered here, where HNLs have hundreds of MeV in mass. With the exception of CHARM and

CHARM-II, the previous experiments operate at energies below the HNL production threshold and are therefore not sensitive to our region of interest. While we could also consider CHARM and CHARM-II, we note that they observe larger neutrino-induced backgrounds thanks to the faster growth of the SM cross section with respect to the dipole one. In addition, the HNLs would be produced with a larger boost factor, and therefore escape more often.

In this work, we will consider three existing measurements of the neutrino-electron cross section by MINERvA [25–27]. The first was performed in the LE configuration of the NuMI beam operating with a forward-horn current (FHC), optimizing the number of neutrinos produced. The last two were performed in the ME configuration, one in FHC and the other in reverse-horn current (RHC) mode, the latter optimizing the number of antineutrinos. The ME RHC measurement, also the most recent, is particularly sensitive due to the smaller antineutrino- and neutrino-induced backgrounds. Unlike the dipole cross section, antineutrino-nucleus weak cross sections are smaller than neutrino-nucleus cross sections.

A. Simulation

As described in Sec. III, we use `LeptonInjector` [28] to simulate $\nu A \rightarrow \mathcal{N}A$ upscattering inside as well as outside the MINERvA detector. A schematic depiction of this process is shown in the top panel of Fig. 2. We include upscattering in the upstream dirt, in the surrounding air, in the nuclear target planes (detailed below), in the plastic scintillator, as well as in the outermost electromagnetic calorimeter. While most of these components are not part of the fiducial volume for the $\nu - e$ analysis, they can significantly contribute to the signal rate due to the displaced decays of the HNLs. For long-lived HNLs, upscattering in the dirt dominates the signal rate, followed by the nuclear target planes, which contain high-density materials like ^{208}Pb and ^{56}Fe . Detailed modeling of the detector geometry and material composition is necessary to correctly predict the contributions from these different upscattering sites. We define each nuclear target to be a hexagonal prism with apothem 92 cm, matching that of the tracker region [61], and z extent given by Table 4 of Ref. [47]. The six nuclear target planes are detailed in Table I. In the bottom panel of Fig. 2, we show the positional distribution of the upscattering rate within two of the MINERvA nuclear targets as simulated using `LeptonInjector` [28]. We also consider upscattering within the electromagnetic calorimeter, defined as a hexagonal prism surrounding the inner detector with 107 cm apothem, and within the steel veto shield ~ 1 m in front of MINERvA [61]. The fiducial volume of MINERvA for the $\nu - e$ analyses is assumed to be approximately the same for both the LE and ME analyses and is defined as a hexagon of 81 cm apothem with ~ 2.8 m z extent inside the plastic scintillator.

After production, we track the HNL's path through the detector and force a decay to occur before the end of

TABLE I. Specifications of the MINERvA nuclear targets as implemented in `LeptonInjector` [28] for this analysis. Each nuclear target is defined as a hexagonal prism with an apothem of 92 cm. Z positions and extents of each nuclear target have been taken from Table 4 of Ref. [47]. The coordinate system is defined such that $z = 0$ corresponds to the front of the MINERvA detector. We have confirmed that the fiducial mass (bounded by an 85 cm apothem hexagon) of each nuclear target sub-component matches the fiducial mass quoted in Table 4 of Ref. [47]. The last column of this table refers to the mass of each nuclear target subcomponent within the 92 cm apothem hexagonal prism.

Target	z -location (cm)	z -extent (cm)	Mass (kg)
1-Fe	13.6	2.567	370
1-Pb	13.6	2.578	317
2-Fe	31.3	2.563	370
2-Pb	31.3	2.581	317
3-Fe	53.4	2.573	197
3-Pb	53.4	2.563	141
3-C	53.4	7.620	194
Water	89.5	18.06	530
4-Pb	125.6	0.795	263
5-Fe	138.9	1.289	186
5-Pb	138.9	1.317	162

the fiducial volume; we then downweight the event by the probability of decaying within the considered region. Each decay produces a photon for which we physically sample a pair-production location. Events that do not pair produce within the fiducial volume are removed. This procedure accounts for events where the HNL decays outside the fiducial volume, but the photon conversion happens inside of it. This effect is important for short-lived HNLs since the rate of HNLs produced in the high-density lead planes can significantly contribute to the signal rate even though they are not contained in the fiducial volume.

The neutrino fluxes for the LE mode have been taken from Ref. [62] and for the ME they have been digitized from Ref. [63]. The total exposures used in the three $\nu - e$ analyses are 3.43×10^{20} POT for LE-FHC, 1.16×10^{21} POT for ME-FHC, and 1.22×10^{21} POT for ME-RHC.

To reduce neutrino-induced backgrounds, MINERvA applies an extensive list of selection cuts. To properly estimate the resulting efficiency of these cuts in our HNL signal, it is important to correctly model the reconstruction of the energy and angle of the single photons. In the absence of a full detector simulation, we proceed to approximate the detector energy and angular resolutions as Gaussian functions. For the energy resolution, we take $\sigma_E/E = 5.9\% / (\sqrt{E/\text{GeV}}) + 3.4\%$ [25], while for the angular resolution, we take an energy-independent angular resolution of $\sigma_\theta = 0.7^\circ$, assumed to be isotropic in the shower's

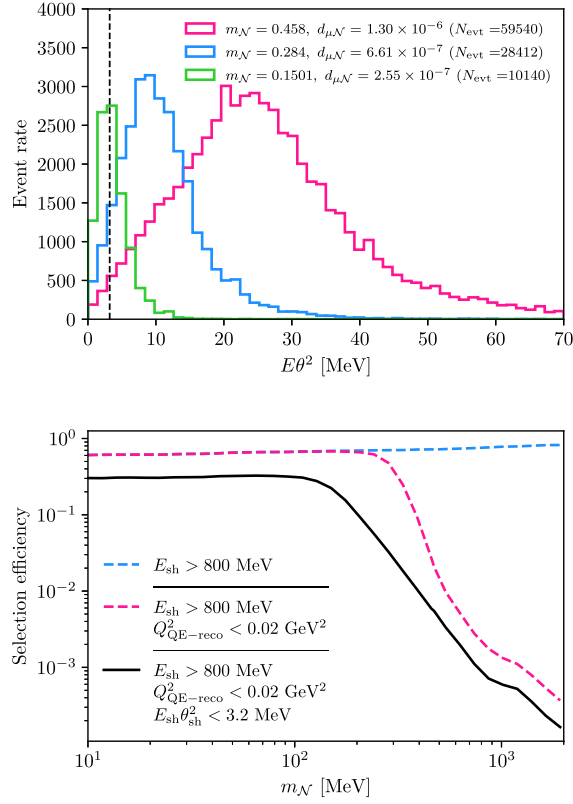


FIG. 8. (Top) The predicted $E_{\text{sh}}\theta_{\text{sh}}^2$ distribution before detector smearing and signal selection for three choices of model parameters at MINERvA. Bottom) The signal selection efficiency of our analysis cuts, excluding the dE/dx cut, as a function of the HNL mass.

azimuthal angle.¹ The angular resolution is implemented by sampling a polar angle $\delta\theta$ from a Gaussian distribution of standard deviation σ_θ , rotating the photon by $\delta\theta$ with respect to its momentum, assigning it an azimuthal angle φ from the uniform distribution $[0, 2\pi]$, and finally rotating the photon back to the laboratory frame by its original polar angle θ_{true} .

B. Event selection

Now we discuss the most important signal selection cuts. The analyses [25,26] make use of a long list of signal selection cuts, designed to suppress as many neutrino-nucleus scattering backgrounds as possible. The most worrisome backgrounds include π^0 production and ν_e CC scattering. The former is particularly important for our radiative decay signal, as it can give rise to coherent single-photonlike signatures. Cuts related to the shower radius and transversal as well as longitudinal profiles are not implemented in our analysis but are expected to have

¹This is only an approximation, as the MINERvA detector is not azimuthally symmetric. Nevertheless, the differences in resolution in the X and Y planes are small [61,64] and neglected here.

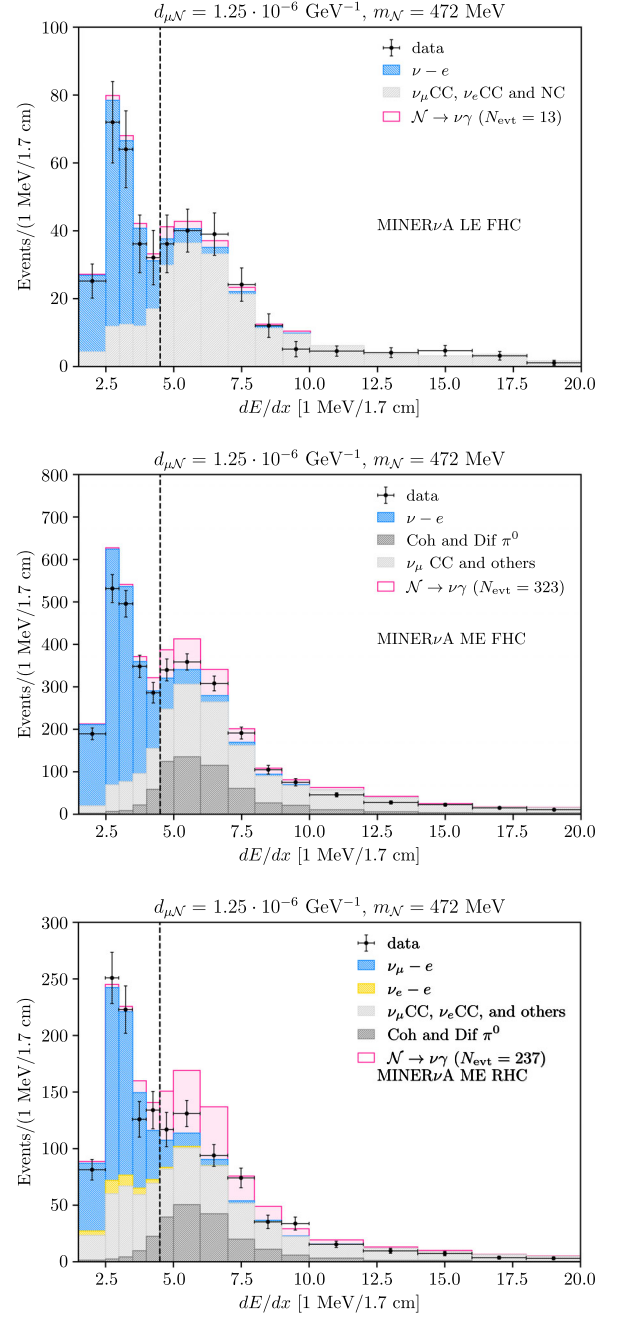


FIG. 9. The dE/dX distribution of selected events for the three MINERvA analyses. From top to bottom: LE FHC, ME FHC, and ME RHC. The latter has the largest sensitivity due to the smaller backgrounds. All distributions are shown post-MINERvA tune, except for ME RHC, where it is shown before tuning.

large acceptance due to our signal being a true single photon (as opposed to two photons from π^0 or from the e^+e^- pairs considered in the new physics model of Ref. [65]). The series of cuts are illustrated in Fig. 8, where we show the acceptance of the cuts as a function of the HNL mass. The selection acceptance is largely independent of the dipole coupling.

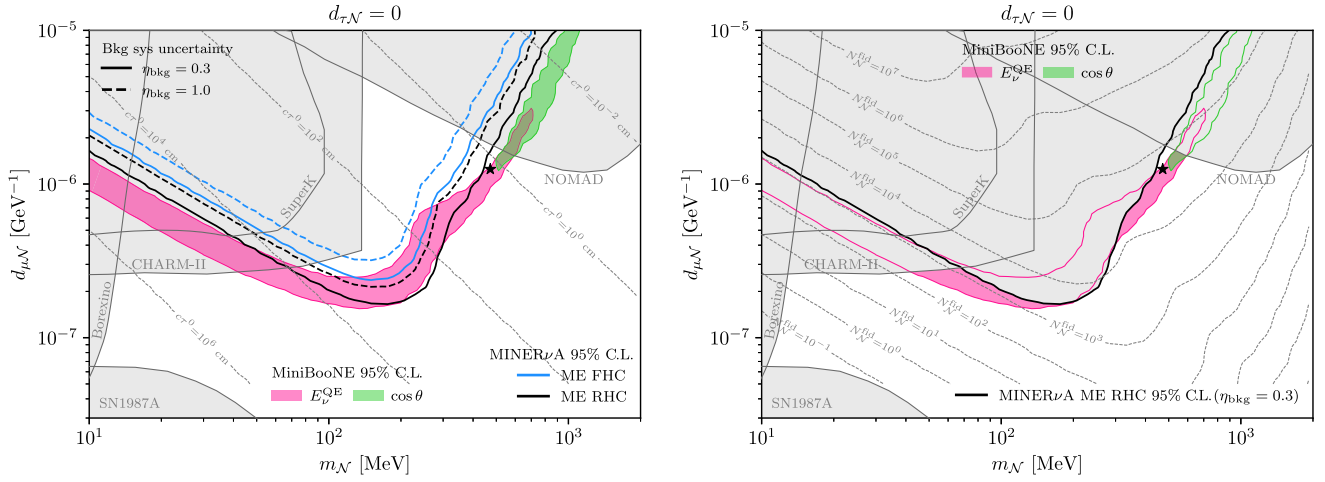


FIG. 10. (Left) MINERvA constraints in the dipole parameter space at 95% C.L. Solid lines show our nominal limits assuming a $\eta_{\text{bkg}} = 30\%$ Gaussian systematic uncertainty on the background normalization, and dashed ones show the constraints assuming an inflated uncertainty of $\eta_{\text{bkg}} = 100\%$ on the background. Regions of preference to explain MiniBooNE in the minimal dipole model are also shown as filled contours at 95% C.L. Right) Contours of constant $N_{\mathcal{N}}^{\text{fid}}$, the total number of new-physics photons that convert inside the fiducial volume, overlaid on top of the same parameter space.

We start with the cut on the reconstructed shower energy, $E_{\text{sh}} > 0.8$ GeV. To suppress ν_e CC backgrounds, a cut on the reconstructed momentum exchange under the hypothesis of neutrino-nucleon quasielastic scattering is also implemented. It is defined as $Q_{\text{QE-reco}}^2 = 2m_n(E_{\nu}^{\text{QE-reco}} - E_{\text{sh}})$, with

$$E_{\nu}^{\text{QE-reco}} = \frac{m_n E_{\text{sh}} - m_e^2/2}{m_n - E_{\text{sh}} + p_e \cos \theta_e}, \quad (6)$$

and the analysis requires $Q_{\text{QE-reco}}^2 < 0.02$ GeV². We note that for $\nu - e$ scattering this cut can be understood by the following relation,

$$Q_{\text{QE-reco}}^2 \simeq (E_{\text{sh}} \theta_{\text{sh}})^2 < \frac{m_e E_{\nu}}{2}, \quad (7)$$

where we dropped higher-order terms in electron mass and θ_{sh} . For neutrino-nucleus upscattering in the forward direction (small $\theta_{\mathcal{N}}$), two-body kinematics for an infinitely heavy nucleus gives $Q_{\text{QE-reco}}^2 \simeq (E_{\mathcal{N}} \theta_{\mathcal{N}})^2 < (E_{\nu} m_A - m_{\mathcal{N}}^2)/2$, which is a much looser constraint. In addition, the decay of \mathcal{N} introduces even more spread in the angular distribution, so we can already expect the cut on $Q_{\text{QE-reco}}^2$ to be very important.

The most stringent cut in the analysis; however, is in $E_{\text{sh}} \theta_{\text{sh}}^2$, required it to be < 3.2 MeV rad². The acceptance of this selection varies from under 10^{-4} at high masses to approximately 30% at the lowest masses. It is larger for Majorana than Dirac HNLs, as Dirac HNLs are more likely to produce backward-going photons. Finally, the last cut we

implement is the cut on the mean dE/dx of first four scintillator planes, requiring $dE/dx > 4.5$ MeV/1.7 cm. Figure 9 shows the dE/dx distributions of the SM background and HNL decay events, as well as a dashed line representing the dE/dx cut. The acceptance of this cut is the least-understood aspect of our analysis as we do not have access to a full detector simulation. The shape of our $dE/dx > 4.5$ MeV/1.7 cm distribution is assumed to be identical to that of the coherent and diffractive π^0 backgrounds shown in Fig. 9. In this approximation, we find this cut has an acceptance of 93% and 91% for FHC and RHC modes, respectively.

C. Results

The final 95% C.L. constraints on the dipole model are shown in Fig. 10. The MiniBooNE regions of preference are also shown for comparison. At the lowest values of $m_{\mathcal{N}}$ where the HNLs are long lived, our constraints are less sensitive to the MiniBooNE best-fit region than at higher masses due to the larger HNL boost factors at MINERvA as well as the smaller fiducial volume when compared to MiniBooNE. For lifetimes longer than $c\tau^0 \sim 100$ cm, the event rate in both experiments is dominated by dirt upscattering. For shorter lifetimes, the event rate is dominated by upscattering within the detector itself.

One can see that the constraints from MINERvA begin to rule out disfavor regions of parameter space preferred by MiniBooNE. However, the strongest MINERvA 2σ C.L. limits presented, which come from the ME RHC measurement and assume 30% uncertainty on the background normalization, do not rule out the intersection of the 2σ C.L. preferred regions from the MiniBooNE E_{ν}^{QE} and $\cos \theta$

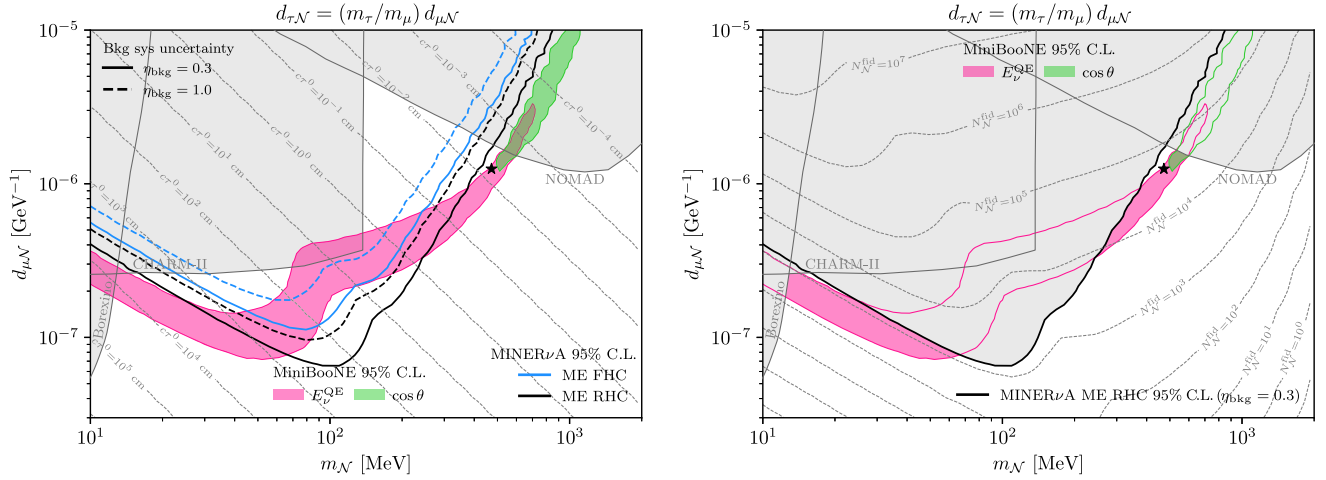


FIG. 11. Same as Fig. 10 but for the case where HNLs have a larger tau-neutrino dipole following an approximate scaling of UV completions of the operator in Eq. (1), $d_{\tau N} = m_\tau/m_\mu \times d_{\mu N}$. The HNL is shorter-lived due to the additional decay $\mathcal{N} \rightarrow \nu_\tau \gamma$.

distributions. This is because the $\mathcal{N} \rightarrow \nu \gamma$ acceptance in the MINERvA ES analysis decreases rapidly for larger HNL masses, as shown in Fig. 8. In the right panel of Fig. 10, we show contours of constant event rate from dipole-coupled HNL decays in MINERvA. A dedicated single-shower analysis improving the acceptance for larger HNL masses would likely be sensitive to the entire region of parameter space preferred by MiniBooNE.

In the left panel of Fig. 10, we also show conservative constraints on this model assuming 100% uncertainty on the background normalization. This is meant to address the large scale factors (up to factors of ~ 2) which have been applied to the high dE/dx backgrounds in the official MINERvA analysis [27]. These scale factors come from a tuning procedure in kinematic sideband regions, a process that could potentially wash out any signal from neutrissimo decays. An optimal analysis would perform a joint fit to both neutrissimo decays and SM high dE/dx backgrounds to derive constraints (and potentially allowed regions) on the neutrissimo model presented here; however, such an analysis is out of the scope of this paper. We also note that the MINERvA analysis does not include single-photon backgrounds such as radiative $\Delta(1232)$ decays and coherent single photons. These components are expected to be small in the energy region of $E_{\text{sh}} > 800$ MeV [66,67], but their inclusion can only make our limits stronger.

We also point out Ref. [68], in which the MINERvA Collaboration investigated an excess in the high dE/dx sideband region of a ν_e charged-current quasielastic scattering sample. Using topological variables related to the shower structure, MINERvA concluded that the excess looked more like coherent or diffractive π^0 production than single photons. This might suggest that the scale factors in Ref. [27] could also be attributed to additional π^0 events. However, the analysis presented here suggests that MINERvA may have unique sensitivity to a

neutrissimo-based explanation of the MiniBooNE excess, thus motivating a more careful separation of one and two photon events in the high dE/dx region of the MINERvA elastic scattering samples.

We also derive constraints considering nonzero $d_{\tau N}$ in Figs. 11 and 12. As expected, this impacts the constraints most significantly at lower HNL masses. For both $d_{\tau N} = (m_\tau/m_\mu)d_{\mu N}$ and $d_{\tau N} = (100 \text{ TeV})^{-1}$ The MINERvA constraints rule out a large chunk in the middle of the region preferred by the MiniBooNE E_ν^{QE} distribution. The constraints do not change for $m_N \gtrsim 200$ MeV, thus the 2σ overlap between between the MiniBooNE E_ν^{QE} and $\cos \theta$ distributions remains valid.

VI. DISCUSSION

A number of other existing and planned neutrino experiments are sensitive to an MeV-scale dipole-coupled HNL. Super-Kamiokande can look for single photon decays from atmospheric neutrinos which upscatter into HNLs within the Earth [69]. Similarly, one can look for single-photon decays from neutrinos that upscatter into HNLs within the Earth and propagate to a large-scale terrestrial detector such as Super-Kamiokande or Borexino [70]. Constraints from these searches are sensitive to longer-lived HNLs, with typical masses $\lesssim 10$ MeV ($\lesssim 100$ MeV) in the solar (atmospheric) case. These constraints no longer apply for the two cases of nonzero $d_{\tau N}$ which we consider in this work, as the HNL lifetime will be too short to reach the detector. Observed neutrino interactions from Supernova 1987A can also be used to derive constraints on the dipole model, as significant upscattering would enhance the stellar cooling rate, decreasing the neutrino flux observed on Earth [18]. These constraints require the HNL to be sufficiently long-lived that it can escape the stellar environment; thus, bounds

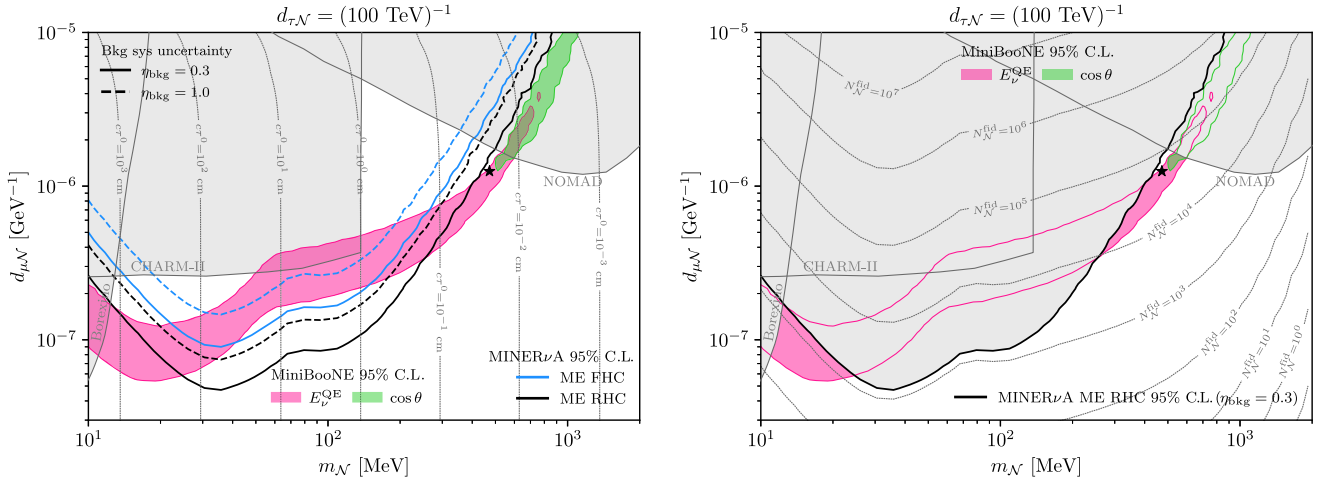


FIG. 12. Same as Fig. 10 but for the case where HNLs have a larger and fixed tau-neutrino dipole, $d_{\tau\mathcal{N}} = (100 \text{ TeV})^{-1}$. The HNL is shorter lived due to the additional decay $\mathcal{N} \rightarrow \nu_\tau \gamma$.

from Supernova 1987A also do not apply when we consider nonzero $d_{\tau\mathcal{N}}$.

We also show constraints derived from the NOMAD search for neutrino-induced single photons, recast as bounds in this parameter space in Ref. [53]. We rescale them, however, by an overall factor of $\sqrt{2}$ so as to reflect the decay rate we derived in Eq. (5). In addition, we note that these limits have been obtained with a much less sophisticated simulation than the ones performed here and that the signal was derived using only events in the preshower detector of NOMAD. In Ref. [53], it is suggested that stronger limits could be obtained by considering upscattering locations beyond the preshower detector and extended detector volumes where the HNL decay could take place.

We also include the limit imposed by CHARM-II as derived in Ref. [71]; however, we note that it was obtained with a simplified procedure. The experimental precision on the total neutrino-electron upscattering ($\nu e \rightarrow \mathcal{N}e$) cross sections. This is a reasonable assumption at low values of $m_{\mathcal{N}}$, but potentially breaks down at values close to the threshold due to differences in kinematics. A similar constraint can be set using the LSND elastic scattering measurement [18], though it is, in general, less sensitive than the CHARM-II measurement. A robust reevaluation of dipole model constraints from these electron scattering measurements is out of the scope of this paper. In addition to scattering on electrons, CHARM-II can provide new limits in the region of interest by considering coherent neutrino-nucleus upscattering followed by HNL decays into single photons [65]. The sideband with large dE/dX and large values of $E\theta^2$ can be used to set limits, as proposed in [65], although we do not expect them to be as sensitive due to larger backgrounds and larger boosts.

One can also derive constraints in the dipole-coupled HNLs from LEP through the $e^+e^- \rightarrow \mathcal{N}\nu_\ell$, which can proceed through either the γ or Z mediators [72]. However, these constraints require a strong enhancement of the mixing between the HNL and SM neutrino; as we consider such a mixing to be negligible in this model, we do not include constraints from LEP in our results.

We now discuss the potential for future constraints on dipole-coupled HNLs from planned measurements. Just like the MINERνA constraint derived in this work, a neutrino elastic scattering measurement from the NOνA experiment would be sensitive to the dipole model [73]. This is especially true of the DUNE experiment, which has the potential to make a high-statistics neutrino-electron scattering measurement [74]. This would be particularly advantageous for the THEIA@DUNE configuration [75] due to its low threshold and large volume. Dedicated searches at neutrino experiments can further improve sensitivity to this model. As discussed above, a MINERνA single-shower analysis without a stringent $E\theta^2$ cut could set much stronger constraints. Experiments which measure CEνNS, such as COHERENT, NUCLEUS, and Coherent CAPTAIN-Mills, would also be sensitive to the dipole model by looking for the coincidence of nuclear recoil from Primakoff upscattering and a single photon from the HNL decay [76]. These experiments would be most sensitive to lower mass HNLs with $m_{\mathcal{N}} \lesssim 10$ MeV due to the lower energy of typical neutrino sources for CEνNS experiments. Existing and upcoming short baseline neutrino experiments, including MicroBooNE and SBND, have the potential to be sensitive to the neutrissimo model presented in this paper through a dedicated search for single photon events [18]. Additionally, neutrino telescopes like IceCube and KM3NeT could perform searches for events with a double-bang topology from the upscattering and decay of the HNL, reaching sensitivities of $d_{\mu\mathcal{N}} \sim 10^{-7} \text{ GeV}^{-1}$ for $m_{\mathcal{N}} \lesssim 1 \text{ GeV}$ [71].

High-energy astrophysical tau neutrino observatories such as TAMBO might also be sensitive to HNL decays from ν_τ upscattering [77]. Projections for DUNE [78,79] estimate that, in the absence of backgrounds, a search for events with a double-bang morphology could reach $d_{\mu\mathcal{N}}$ values as low as $\mathcal{O}(10^{-8} \text{ GeV}^{-1})$. Finally, nuclear emulsion and liquid argon detectors at a future LHC Forward Physics Facility will also be sensitive to transition magnetic moments between HNLs and SM neutrinos [80].

VII. CONCLUSION

In this work, we have explored a mixed model comprising an eV-scale sterile neutrino and an MeV-scale dipole-coupled HNL. The former facilitates oscillations at short baselines, while the latter introduces the interactions shown in Fig. 1. The dipole-coupled HNL provides an alternative explanation of the MiniBooNE excess to the eV-scale sterile neutrino. Thus one can remove MiniBooNE from global $3 + 1$ fits, reducing tension between appearance and disappearance experiments while retaining an explanation of the LSND anomaly [23]. We take the result of the MiniBooNE-less $3 + 1$ global fit as the oscillation contribution to the MiniBooNE excess and attribute the remaining excess to decays of the dipole-coupled HNL. We find that spectral fits to the E_D^{QE} and $\cos\theta$ distributions prefer different regions of parameter space in general, though solutions exist which are compatible with both distributions at the 2σ confidence level.

We have also derived constraints on the dipole-coupled HNL model using a MINERvA neutrino-electron elastic scattering measurements [26,27,61]. We find that the most sensitive $\nu - e$ scattering constraints are those obtained with the NuMI medium-energy mode in antineutrino-enhanced beam configuration. The constraints from antineutrino-mode are especially strong due to a reduction in backgrounds at high dE/dx , where we expect HNL decays to contribute. As shown in Fig. 10, MINERvA can exclude large regions of parameter space preferred by MiniBooNE, but it does not fully exclude it. There are still allowed MiniBooNE regions at the 2σ confidence level. The first is at small $m_{\mathcal{N}}$ values, where HNLs are long-lived and MINERvA's small fiducial volume and larger energies reduce the sensitivity. The second is at larger HNL masses, where the stringent $E\theta^2$ cuts reject most new physics events where decay photons tend to have larger $E\theta^2$. We note that a dedicated search at MINERvA using the same fiducial volume could significantly improve the signal efficiency in this large-mass region, and would likely have much better sensitivity, and potentially probe the entire MiniBooNE-preferred region, as shown in the right panel of Fig. 10. Nevertheless, as it stands, this mixed model of oscillations and decay is not dead yet.

ACKNOWLEDGMENTS

M. H. S. is supported by NSF Grant No. PHY-1707971. NSF Grant No. PHY-1801996 supported C. A. A., J. M. C., A. D., and N. W. K. for this work. Additionally, C. A. A. is supported by the Faculty of Arts and Sciences of Harvard University and the Alfred P. Sloan Foundation. N. W. K. is supported by the NSF Graduate Research Fellowship under Grant No. 1745302. M. A. U. is supported by the Department of Physics at the University of Cambridge and S. V. is supported by the STFC. M. H. was supported by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research, Innovation, and Science.

APPENDIX: UPSCATTERING CROSS SECTION

The cross section for \mathcal{N} production by neutrino upscattering, $\nu_\alpha + A \rightarrow \mathcal{N} + A$, has been computed several times in the literature [41,81,82]. We note, however, that all expressions we could find do not take into account the polarization of the outgoing HNL. This effect is only important when $m_{\mathcal{N}}/E$ becomes appreciably large and we find it to be a marginal effect in our calculations.

In the massless limit, a beam of left-handed polarized neutrinos will always upscatter to right-handed polarized HNLs, assuming the process takes place purely via the transition magnetic moment. This follows from the chiral structure of the vertex, $\bar{\nu}_L \sigma^{\mu\nu} \mathcal{N}_R$. However, in the massive case, spin and helicity are not equivalent, and both helicity states of \mathcal{N} can be produced. The helicity-flipping channel, $\nu_h \rightarrow \mathcal{N}_{-h}$, typically dominates, while the helicity-conserving case, $\nu_h \rightarrow \mathcal{N}_h$, will be suppressed by powers of $m_{\mathcal{N}}/E$, vanishing in the massless limit. Here $h = \pm 1$ denotes the particle's helicity and E is the typical energy scale of the scattering process.

We have calculated both terms using the `DarkNews` code [83], and show our results in Fig. 14. We show a comparison of the upscattering cross section on Carbon-12 for a few choices of HNL masses for both coherent and proton-elastic scattering regimes. Scattering on neutrons proceeds only via the neutron magnetic moment and is much smaller. We only include the proton-elastic contribution for MiniBooNE, where the proton would be invisible. This is a conservative approach when deriving the MINERvA limits.

We also show the ratio between helicity-flipping and helicity-conserving upscattering events inside the MiniBooNE detector as a function of $m_{\mathcal{N}}$ in Fig. 15. The helicity-conserving part is a small correction, except at the very largest HNL masses, where the rate is significantly smaller due to the large energy threshold for upscattering.

Several model-independent nuclear form factor parametrizations can be found in the literature. One of them is




FIG. 13. Comparison of the different nuclear form factors for ^{12}C commonly used in the literature. The dipole form factor (dashed blue) significantly overestimates the form factor at large values of the momentum exchange Q . In this paper, we use the Fourier-Bessel parametrization (solid green) for nuclei for which nuclear data is available, otherwise, we implement the Fermi-symmetrized (FS) Woods-Saxon parametrization (dotted yellow).




FIG. 14. Comparison between helicity-flipping and helicity-conserving cross sections for coherent neutrino upscattering on Carbon (left) and free protons (right) for multiple values of the HNL mass $m_{\mathcal{N}}$.




FIG. 15. The ratio between helicity-conserving and helicity-flipping upscattering events on Carbon at MiniBooNE.

the Fourier-Bessel parametrization, which models the charge density in the nucleus as a series of Bessel functions with a radial cutoff of R .

$$F^{\text{FB}}(Q) = N \times \frac{\sin(QR)}{QR} \sum_n \frac{(-1)^n a_n}{n^2 \pi^2 - Q^2}, \quad (\text{A1})$$

where $1/N = \sum_n (-1)^n a_n / n^2 \pi^2$ is a normalization factor, ensuring $F(0) = 1$. The coefficients a_n can be obtained from experimental data, which is available for a series of common nuclei [48–50]. We make use of the machine-readable files provided by Ref. [51].

For nuclei where the nuclear data cannot be found, we implement a Fermi-symmetrized Woods-Saxon form factor,

$$F^{\text{FS-WS}}(Q) = \frac{3\pi a}{r_0^2 + \pi^2 a^2} \frac{a\pi \cotanh(\pi Q a) \sin(Q r_0) - r_0 \cos(Q r_0)}{Q r_0 \sinh(\pi Q a)}, \quad (\text{A2})$$

where $a = 0.523$ fm, $r_0 = 1.03 \times A^{1/3}$ fm. These form factors correctly describe the finite nuclear radius and lead to a strong suppression of coherent scattering for $Q \gtrsim 200$ GeV. They should be contrasted with the simpler dipole parametrization

$$F^{\text{Dip}}(Q) = \frac{1}{1 + \frac{Q^2}{M_D^2}}. \quad (\text{A3})$$

used in Ref. [23], with $M_D = 1.18 + 0.83 * A^{1/3}$, and the Helmz form factor

$$F^{\text{Helmz}}(Q) = \frac{3|j_1(QR)|}{QR} e^{-Q^2 s^2/2}, \quad (\text{A4})$$

with $a = 0.523$ fm, $s = 0.9$ fm, and $R = 3.9$ fm.

We provide a comparison of the aforementioned nuclear form factors in Fig. 13. It is evident that the dipole parametrization overestimates the cross section at large values of momentum exchange Q . The more sophisticated form factors used in this work produce more forward angular distributions at MiniBooNE than what was found in the previous study of Ref. [23].

-
- [1] Q. R. Ahmad *et al.* (SNO Collaboration), *Phys. Rev. Lett.* **89**, 011302 (2002).
- [2] Y. Ashie *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. D* **71**, 112005 (2005).
- [3] P. F. de Salas, D. V. Forero, S. Gariazzo, P. Martínez-Miravé, O. Mena, C. A. Ternes, M. Tórtola, and J. W. F. Valle, *J. High Energy Phys.* **02** (2021) 071.
- [4] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou, *J. High Energy Phys.* **09** (2020) 178.
- [5] F. Capozzi, E. Di Valentino, E. Lisi, A. Marrone, A. Melchiorri, and A. Palazzo, *Phys. Rev. D* **104**, 083031 (2021).
- [6] A. A. Aguilar-Arevalo *et al.* (LSND Collaboration), *Phys. Rev. D* **64**, 112007 (2001).
- [7] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Phys. Rev. D* **103**, 052002 (2021).
- [8] M. A. Acero *et al.*, Report No. FERMILAB-PUB-22-318-ND-SCD-T, 2022.
- [9] W. Loinaz, N. Okamura, S. Rayyan, T. Takeuchi, and L. C. R. Wijewardhana, *Phys. Rev. D* **70**, 113004 (2004).
- [10] T. Adams *et al.* (NuSOnG Collaboration), *Int. J. Mod. Phys. A* **24**, 671 (2009).
- [11] S. N. Gninenko, *Phys. Rev. Lett.* **103**, 241802 (2009).
- [12] S. N. Gninenko, *Phys. Rev. D* **83**, 015015 (2011).
- [13] C. Dib, J. C. Helo, S. Kovalenko, and I. Schmidt, *Phys. Rev. D* **84**, 071301 (2011).
- [14] S. N. Gninenko, *Phys. Lett. B* **710**, 86 (2012).
- [15] M. Masip, P. Masjuan, and D. Meloni, *J. High Energy Phys.* **01** (2013) 106.
- [16] A. Radionov, *Phys. Rev. D* **88**, 015016 (2013).
- [17] P. Ballett, S. Pascoli, and M. Ross-Lonergan, *J. High Energy Phys.* **04** (2017) 102.
- [18] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, *Phys. Rev. D* **98**, 115015 (2018).
- [19] A. B. Balantekin, A. de Gouvêa, and B. Kayser, *Phys. Lett. B* **789**, 488 (2019).
- [20] S. Balaji, M. Ramirez-Quezada, and Y.-L. Zhou, *J. High Energy Phys.* **04** (2020) 178.
- [21] S. Balaji, M. Ramirez-Quezada, and Y.-L. Zhou, *J. High Energy Phys.* **12** (2020) 090.
- [22] O. Fischer, A. Hernández-Cabezudo, and T. Schwetz, *Phys. Rev. D* **101**, 075045 (2020).
- [23] S. Vergani, N. W. Kamp, A. Diaz, C. A. Argüelles, J. M. Conrad, M. H. Shaevitz, and M. A. Uchida, *Phys. Rev. D* **104**, 095005 (2021).
- [24] L. Alvarez-Ruso and E. Saul-Sala, *Eur. Phys. J. Special Topics* **230**, 4373 (2021).
- [25] J. Park *et al.* (MINERvA Collaboration), *Phys. Rev. D* **93**, 112007 (2016).
- [26] E. Valencia *et al.* (MINERvA Collaboration), *Phys. Rev. D* **100**, 092001 (2019).
- [27] L. Zazueta *et al.* (MINERvA Collaboration), *Phys. Rev. D* **107**, 012001 (2023).
- [28] R. Abbasi *et al.* (IceCube Collaboration), *Comput. Phys. Commun.* **266**, 108018 (2021).
- [29] K. Fujikawa and R. Shrock, *Phys. Rev. Lett.* **45**, 963 (1980).
- [30] P. B. Pal and L. Wolfenstein, *Phys. Rev. D* **25**, 766 (1982).
- [31] R. E. Shrock, *Nucl. Phys.* **B206**, 359 (1982).
- [32] C. A. Argüelles, N. Foppiani, and M. Hostert, *Phys. Rev. D* **105**, 095006 (2022).
- [33] M. B. Voloshin, *Sov. J. Nucl. Phys.* **48**, 512 (1988).
- [34] R. Barbieri and R. N. Mohapatra, *Phys. Lett. B* **218**, 225 (1989).
- [35] K. S. Babu and R. N. Mohapatra, *Phys. Rev. Lett.* **64**, 1705 (1990).
- [36] K. S. Babu and R. N. Mohapatra, *Phys. Rev. Lett.* **63**, 228 (1989).
- [37] M. Leurer and N. Marcus, *Phys. Lett. B* **237**, 81 (1990).
- [38] M. Lindner, B. Radovčić, and J. Welter, *J. High Energy Phys.* **07** (2017) 139.
- [39] K. S. Babu, S. Jana, and M. Lindner, *J. High Energy Phys.* **10** (2020) 040.

- [40] A. A. Aguilar-Arevalo *et al.* (MiniBooNE), *Phys. Rev. Lett.* **129**, 201801 (2022).
- [41] V. Brdar, A. Greljo, J. Kopp, and T. Opferkuch, *J. Cosmol. Astropart. Phys.* **01** (2021) 039.
- [42] V. Brdar, O. Fischer, and A. Y. Smirnov, *Phys. Rev. D* **103**, 075008 (2021).
- [43] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), [arXiv:1207.4809](https://arxiv.org/abs/1207.4809).
- [44] MiniBooNE Collaboration (private communication).
- [45] M. H. Shaevitz (MiniBooNE Collaboration), *J. Phys. Conf. Ser.* **120**, 052003 (2008).
- [46] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **599**, 28 (2009).
- [47] L. Aliaga *et al.* (MINERvA Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **743**, 130 (2014).
- [48] G. Fricke, C. Bernhardt, K. Heilig, L. A. Schaller, L. Schellenberg, E. B. Shera, and C. W. de Jager, *At. Data Nucl. Data Tables* **60**, 177 (1995).
- [49] H. De Vries, C. W. De Jager, and C. De Vries, *At. Data Nucl. Data Tables* **36**, 495 (1987).
- [50] C. W. De Jager, H. De Vries, and C. De Vries, *At. Data Nucl. Data Tables* **14**, 479 (1974); **16**, 580(E) (1975).
- [51] <http://discovery.phys.virginia.edu/research/groups/ncd/index.html>.
- [52] S. S. Wilks, *Ann. Math. Stat.* **9**, 60 (1938).
- [53] S. N. Gninenko and N. V. Krasnikov, *Phys. Lett. B* **450**, 165 (1999).
- [54] J. Dorenbosch *et al.* (CHARM Collaboration), *Z. Phys. C* **41**, 567 (1989); **51**, 142 (1991).
- [55] P. Vilain *et al.* (CHARM-II Collaboration), *Phys. Lett. B* **335**, 246 (1994).
- [56] L. B. Auerbach *et al.* (LSND Collaboration), *Phys. Rev. D* **63**, 112001 (2001).
- [57] M. Deniz *et al.* (TEXONO Collaboration), *Phys. Rev. D* **81**, 072001 (2010).
- [58] G. Bellini *et al.*, *Phys. Rev. Lett.* **107**, 141302 (2011).
- [59] M. Agostini *et al.* (Borexino Collaboration), *Phys. Rev. D* **100**, 082004 (2019).
- [60] S. Fukuda *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. Lett.* **86**, 5651 (2001).
- [61] J. Park, Ph.D. thesis, University of Rochester, 2013.
- [62] L. Aliaga *et al.* (MINERvA Collaboration), *Phys. Rev. D* **94**, 092005 (2016); **95**, 039903(A) (2017).
- [63] A. Bashyal, Ph.D. thesis, Oregon State University, 2021.
- [64] E. Valencia-Rodriguez, Ph.D. thesis, Guanajuato University, 2016.
- [65] C. A. Argüelles, M. Hostert, and Y.-D. Tsai, *Phys. Rev. Lett.* **123**, 261801 (2019).
- [66] D. Rein and L. M. Sehgal, *Phys. Lett.* **104B**, 394 (1981); **106B**, 513(E) (1981).
- [67] E. Wang, L. Alvarez-Ruso, and J. Nieves, *Phys. Lett. B* **740**, 16 (2015).
- [68] J. Wolcott *et al.* (MINERvA Collaboration), *Phys. Rev. Lett.* **117**, 111801 (2016).
- [69] R. A. Gustafson, R. Plestid, and I. M. Shoemaker, *Phys. Rev. D* **106**, 095037 (2022).
- [70] R. Plestid, *Phys. Rev. D* **104**, 075027 (2021).
- [71] P. Coloma, P. A. N. Machado, I. Martinez-Soler, and I. M. Shoemaker, *Phys. Rev. Lett.* **119**, 201804 (2017).
- [72] F. Delgado, L. Duarte, J. Jones-Perez, C. Manrique-Chavil, and S. Peña, *J. High Energy Phys.* **09** (2022) 079.
- [73] J. Bian, in *Proceeding for the APS DPF 2017 conference in Batavia, IL* (2017), <https://inspirehep.net/conferences/1430017?ui-citation-summary=1430017?ui-citation-summary=true>.
- [74] C. M. Marshall, K. S. McFarland, and C. Wilkinson, *Phys. Rev. D* **101**, 032002 (2020).
- [75] M. Askins *et al.* (Theia Collaboration), *Eur. Phys. J. C* **80**, 416 (2020).
- [76] P. D. Bolton, F. F. Deppisch, K. Fridell, J. Harz, C. Hati, and S. Kulkarni, *Phys. Rev. D* **106**, 035036 (2022).
- [77] A. Romero-Wolf *et al.*, in *Latin American Strategy Forum for Research Infrastructure* (2020), <https://inspirehep.net/conferences/1726700?ui-citation-summary=true>.
- [78] M. Atkinson, P. Coloma, I. Martinez-Soler, N. Rocco, and I. M. Shoemaker, *J. High Energy Phys.* **04** (2022) 174.
- [79] T. Schwetz, A. Zhou, and J.-Y. Zhu, *J. High Energy Phys.* **07** (2021) 200.
- [80] A. Ismail, S. Jana, and R. M. Abraham, *Phys. Rev. D* **105**, 055008 (2022).
- [81] P. Vogel and J. Engel, *Phys. Rev. D* **39**, 3378 (1989).
- [82] R. Harnik, J. Kopp, and P. A. N. Machado, *J. Cosmol. Astropart. Phys.* **07** (2012) 026.
- [83] A. Abdullahi, J. Hoefken, M. Hostert, D. Massaro, and S. Pascoli, *DarkNews*, <https://github.com/mhostert/DarkNews> (2019).

Chapter 7

The Coherent CAPTAIN-Mills Experiment

We now turn to the Coherent CAPTAIN-Mills (CCM) experiment at Los Alamos National Laboratory (LANL). CCM uses a 10-ton cylindrical liquid argon detector to observe the scintillation and Cherenkov light generated by the interactions of particles produced at the Lujan beam dump facility [30]. CCM was originally designed to detect ν_μ produced by π^+ decay-at-rest via coherent elastic neutrino-nucleus scattering (CE ν NS) [269]. Given the baseline and neutrino energy accessible to the CCM detector, any deficit in this rate would be smoking-gun evidence for ν_μ disappearance via a sterile neutrino [30]. In order to realize this physics goal, CCM was built to maximize photon yield and thus minimize the detection energy threshold. This setup makes CCM an ideal detector to search for potential light dark matter (DM) particles produced in the Lujan beam dump, including $U(1)'$ vector-portal DM [30], leptophobic vector-portal DM [270], and axion-like particles (ALPs) [271]. CCM120, a 120-PMT prototype detector, completed a six-week engineering run in Fall 2019, which was used to set constraints on these models. The full 200-PMT CCM200 detector is funded for a three-year physics run at the Lujan facility, which began in Summer 2022. Results from this dataset will set strong constraints on the DM models listed above as well as a variety of additional beyond-the-Standard-Model (BSM) scenarios.

In this chapter, we discuss the projects within CCM in which I was involved

during my time as a graduate student. These include the fabrication of veto PMT assemblies for CCM200, the investigation of Cherenkov light detection as a powerful background rejection technique in CCM, and the evaluation of CCM’s sensitivity to the neutrissimo model from chapter 6.

7.1 The CCM Beamline and Detector

The CCM detector operates at the Lujan beam dump facility of the Los Alamos Neutron Science Center (LANSCE) [272,273]. At LANSCE, protons are accelerated to 800 MeV in the proton storage ring and bunched before impinging vertically downward on the Lujan tungsten target at a rate of 20 Hz. The bunches contain $\sim 3.1 \times 10^{13}$ protons each and are delivered in a 280 ns triangular pulse. The beam spills occur at a rate of 20 Hz, corresponding to a duty factor of $\sim 5 \times 10^{-6}$. This low duty factor is essential for isolating prompt signal events from background, setting the Lujan facility apart from similar beam dump sources. A cascade of particles is produced when protons hit the tungsten target. The target is housed within a cylindrical target-moderator-reflector-shield (TMRS) which provides immediate shielding around the tungsten disks as well as moderators to control neutron output [28], as shown in figure 7-1. The TMRS itself is surrounded by an additional 4 m of steel shielding. Though the primary objective of the Lujan beam dump is neutron production, proton collisions also create pions, electrons, photons, and potential DM particles. Neutrinos are created via the $\pi^+ \rightarrow \nu_\mu(\mu^+ \rightarrow \bar{\nu}_\mu \nu_e e^+)$ decay chain, as π^- dominantly capture on nuclei instead of decaying. The total neutrino flux produced by the Lujan source for each neutrino species in this decay chain is $4.74 \times 10^5 \nu/\text{cm}^2/s$ at the CCM120 location (20 m from the tungsten target). The π^+ come to rest in the target before decaying, leading to the neutrino energy distributions shown in figure 7-2a. Figure 7-2b shows the timing profile of different particles produced in the beam dump at the CCM120 location. This timing distribution is essential for separating ν_μ from π^+ decay-at-rest and speed-of-light DM particles in the prompt region from delayed non-relativistic neutron backgrounds. This is the main strategy behind the CCM physics analyses.

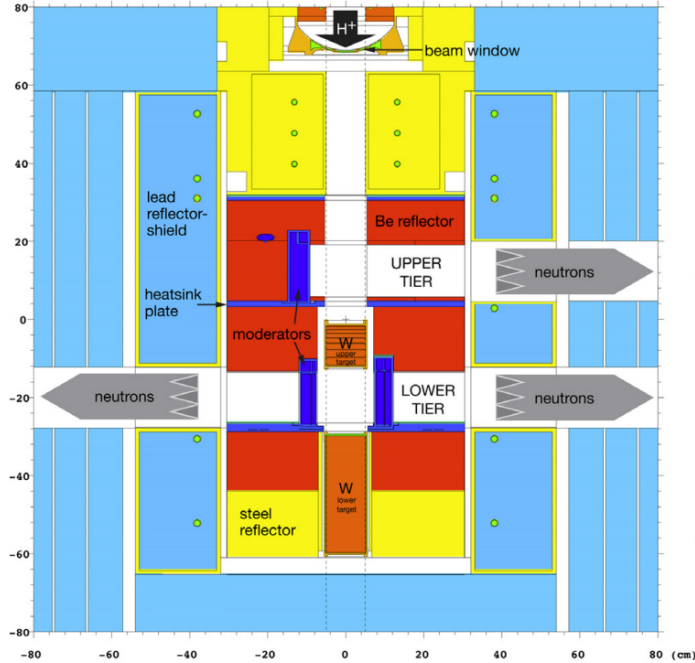


Figure 7-1: Schematic depiction of the Lujan TMRS Figure from Ref. [28].




Figure 7-2: Figure 7-3a, from Ref. [29], shows the energy distribution of π^+ decay-at-rest neutrinos from the Lujan beam dump source. Figure 7-2b, from Ref. [30], shows the timing distribution of particles produced in the Lujan beam dump source after traveling through the TMRS.

The CCM detector is a 10-ton cylindrical volume of liquid argon 2.58 m diameter and 2.25 m in height. CCM repurposes the cryostat used by the Cryogenic Apparatus for Precision Tests of Argon Interactions with Neutrinos (CAPTAIN) experiment [274]. In order to take advantage of the timing information discussed above, CCM relies on the fast scintillation light produced when particles interact in the de-

tector, which is emitted on a characteristic time scale of $\mathcal{O}(ns)$. Scintillation light in liquid argon was discussed in section 3.1.3 within the context of the MicroBooNE experiment. In order to collect this scintillation light, the CCM200 detector is instrumented with 200 8-inch Hamamatsu R5912-mod2 PMTs, the same used in the MicroBooNE light collection system described in section 3.1.3. A schematic drawing of this detector is shown in figure 7-3a. 80 PMTs are arranged uniformly on the top and bottom of the detector, while 120 PMTs are arranged in five rows around the barrel. The CCM120 engineering run only included the PMTs around the barrel. 80% of the PMTs are coated with TPB, which shifts the 128 nm argon wavelength to the visible regime where it can be detected by the PMTs. Additionally, the walls of the detector are covered in reflective foils that are evaporatively coated with TPB. The remaining 20% of PMTs are not coated with TPB and are thus not sensitive to the direct scintillation light produced in the detector. They are only able to observe photons after they are re-emitted by a TPB-coated surface somewhere else in the detector, either on another PMT or on the reflective foils. The presence of uncoated PMTs breaks the degeneracy of the detector response to 128 nm light, helping determine the TPB properties during the calibration process. This point will be important later in section 7.2. Figure 7-3b shows an image of the interior PMTs of the CCM200 detector, in which the dull (reflective) hemispheres correspond to coated (uncoated) PMTs. The voltage output from each PMTs is carried through warm cables to a series of CAEN VX1730 boards, which digitize the PMT signals at a rate of 500 MHz. While the beam is running, the digitizers are triggered at a rate of 22.2 Hz, which includes the beam trigger (20 Hz), random strobe trigger (1.1 Hz), and light-emitting-diode (LED) trigger (1.1 Hz). For each trigger, 16 μs of waveform data on each PMT are saved.

CCM200 also houses an additional 40 veto PMTs outside of the fiducial volume to identify and reject events that originate outside of the fiducial volume. These veto PMTs observe light The CCM120 detector used only 28 veto PMTs, of which 5 were the same 8-inch Hamamatsu R5912-mod2 PMTs used in the inner volume and 23 were 1-inch Hamamatsu PMTs. An additional 20 1-inch veto PMTs were assembled




Figure 7-3: Figure 7-3a shows a schematic 3D rendering of the CCM200 detector. Figure 7-3b shows an image of the interior of the CCM200 detector.

and tested at MIT in Summer 2021. Each PMT was assembled into the mounts shown in figure 7-4, which included an acrylic window in front of the photocathode and a custom base circuit for reading out the PMT signal. The acrylic windows were coated with a mixture comprising 1 g of TPB, 1 g of polystyrene, and 50 mL of (very pungent) toluene. The PMTs were tested in a light-tight box with an LED, as shown in figure 7-5. The LED was pulsed via a function generator at 1 V and 2 V, leading to the average veto PMT response shown in section 7.1. The overshoot observed in these waveforms came largely from the custom breadboard “splitter” circuit assembled for this study, which separated transient PMT signals from high voltage. As the official splitter circuit used by CCM is more robust to ringing, this response was deemed sufficient for CCM200, and the 1-inch veto PMTs were shipped back to LANL.

CCM uses a variety of shielding between the target and detector to slow down and absorb beam-related neutrons, the most prominent beam-related background at the Lujan facility. This reduces the overall background rate and also separates beam-related neutrons in time from prompt BSM signals. Figure 7-2b shows the timing distributions of neutrons, π^+ decay-at-rest neutrinos, and prompt DM from the Lujan source after the 4 m of steel shielding surrounding the TMRS. The neutrons are




Figure 7-4: Two of the veto PMT assemblies constructed at MIT, including the 1-inch PMT, base circuit board, and TPB-coated acrylic window.





Figure 7-5: Section 7.1 shows one of the veto PMTs across from the LED in the light-tight box. Section 7.1 shows the average response of the 20 veto PMTs to 1 V (top) and 2 V (bottom) LED pulses.

pushed back further in time using additional shielding concrete, steel, and borated polyethylene shielding. For the CCM120 run, 6 ft of concrete and 16 in. of steel were placed immediately outside of the steel surrounding the TMRS. A 2 m-thick

wall of steel was constructed in ER2, the room housing the CCM detector. Three additional walls of concrete bricks covered with borated polyethylene sheets were constructed around the front and sides of the CCM detector. The borated polyethylene sheets were included to absorb neutrons that had been sufficiently slowed down by the other shielding. An Eljen EJ-301 scintillation detector was deployed in a flight path neighboring CCM to observe the time of the prompt gamma flash from the Lujan source. Figure 7-6 shows the timing distribution of these photons compared to that of neutrons measured in the CCM120 detector. The discrepancy between these distributions was used to determine a 190 ns background-free region of interest (ROI) in which to look for prompt signals from neutrinos or BSM particles. The steady-state background expectation in the ROI is taken from a beam-out-of-time measurement before the ROI, which is about 22 times larger than the ROI itself [30].

A number of updates were made to the shielding for CCM200. A new larger wall of concrete, steel, and borated polyethylene was constructed in ER1, and more steel and concrete shielding was placed in ER2. Additionally, a steel and polyethylene roof was constructed over CCM to suppress backgrounds entering the detector from above. Preliminary estimates suggest the shielding upgrades have reduced backgrounds in CCM200 by a factor of 7 compared with CCM120 [275].

CCM uses three different methods to calibrate the detector. The first of these consists of two blue LEDs, which are used to determine the single photo-electron (p.e.) response of each PMT, as measured in units of analog-to-digital-conversion (ADC) counts. The PMTs are gain-matched in order to make the p.e.-to-ADC values as similar as possible between them. These LEDs are operated continuously while taking beam data in order to monitor changes in the single p.e. rate over time. CCM also uses ^{57}Co and ^{22}Na radioactive sources, which emit 126 keV and 2.2 MeV photons, respectively, to determine the energy scale of the detector. This is measured in units of p.e./MeVee (electromagnetic equivalent)—a quenching factor must be considered to determine the energy scale of the detector response to nuclear recoils [276, 277]. The sodium peak, which provided the more robust energy scale measurement, set the energy scale of CCM120 to 15.1 ± 4.0 p.e./MeVee. Finally, a custom diode-

pumped laser is used to determine the PMT response to 213 nm and 532 nm light. The ability to operate at two different wavelengths is important for disentangling the TPB response in CCM. Data from the radioactive sources and the laser were compared with output from a `Geant4` simulation of CCM120 in order to tune the optical model of the detector, as described in Ref. [275].

During the operation of CCM120, the effective PMT quantum efficiency was about a factor of two lower than the expectation from Hamamatsu. This was likely related to the oxygen contamination, which was measured at the 2 ppm level in the CCM120 liquid argon. Such contamination can lead to a reduction of the scintillation light produced in liquid argon by about a factor of two [164], similar to the apparent reduction in quantum efficiency. In response to this, CCM200 will use a dedicated oxygen filtration system adapted from the Mini-CAPTAIN experiment [278] to keep the oxygen level below 100 ppb. Filtration is expected to increase the scintillation light output by a factor of around three [275].




Figure 7-6: The timing distribution of photons from the Lujan source (solid black line) compared with that of neutrons measured in CCM120 (dashed red line). Figure from Ref. [30].

7.2 Cherenkov Light Reconstruction

CCM relies predominantly on scintillation light to reconstruct particle interactions in the detector. This is because liquid argon is a powerful scintillation medium, as described in section 3.1.3. That being said, the identification of Cherenkov radiation in CCM is a promising potential upgrade. Cherenkov radiation refers to the light emitted by a charged particle as it travels faster than the speed of light in a given medium [279]. The particle polarizes surrounding atoms as it moves through the material, creating a cloud of electric dipoles. If $\beta > c/n$ (where n is the wavelength-dependent index of refraction), the dipoles are not arranged symmetrically around the particle, resulting in a non-vanishing overall dipole moment. The time variation of this dipole moment results in Cherenkov radiation. The light is emitted in a cone with an opening angle of $\cos \theta_C = 1/n\beta$. The number of photons emitted per unit length is given by [279]

$$\frac{dN}{dx} = 2\pi\alpha z^2 \int_{\lambda_1}^{\lambda_2} \left(1 - \frac{1}{(n(\lambda)\beta)^2}\right) \frac{d\lambda}{\lambda^2}, \quad (7.1)$$

where z is the electric charge of the particle in units of the electron charge e , and λ_1 and λ_2 represent the photon wavelength range under consideration. From equation (7.1), one can see that the intensity of Cherenkov radiation peaks at lower wavelengths.

For liquid argon, $n(\lambda = 128 \text{ nm}) \sim 1.5$ and decreases for larger wavelengths [167]. This means that only electrons satisfy the $\beta > c/n$ requirement at typical CCM energies of $\mathcal{O}(10 \text{ MeV})$. Though they are neutral, photons in this energy regime will produce Cherenkov light after undergoing pair production to an e^+e^- final state. Crossing cosmic muons will also produce Cherenkov light, though such events are rare in the $1.6 \mu\text{s}$ beam window and can be tagged by the veto region. Importantly, the most prominent backgrounds in CCM–beam-related neutrons and low-energy photons from radioactive impurities–will not produce Cherenkov light. In contrast, signals from ALPs [271] and other BSM scenarios of interest to CCM [145] will produce Cherenkov light via higher-energy electrons and photons in the final state. Thus,

Cherenkov light identification is a promising method for rejecting backgrounds in CCM and maximizing sensitivity to new physics. The directional nature of Cherenkov light will also be helpful in requiring particles to originate from the tungsten target, which is not possible using only isotropic scintillation light.

In order to determine the number of Cherenkov photons generated per unit length traveled by a charged particle, one must perform the integration equation (7.1). To do this, we use the parameterization of the liquid argon index of refraction given in Ref. [167]. We fix $\lambda_2 = 700$ nm, as this is the upper bound on the detectable wavelength range of the CCM PMTs. Figure 7-7a shows a numerical integration of equation (7.1) for an electron ($z = 1$) as a function of λ_1 . One can see that sensitivity to lower wavelengths greatly increases the number of photons one can detect. This is possible in CCM thanks to the use of TPB, which can convert the UV Cherenkov light into the visible range, greatly increasing the number of Cherenkov photons detected by each PMT. The cone angle as a function of the electron kinetic energy and photon wavelength is shown in figure 7-7b, which is $\cos \theta_C \approx 0.8$ over most of the parameter space. This figure also indicates the threshold for Cherenkov emission; over most photon wavelengths, Cherenkov light is produced for electron kinetic energies above $T_{e^-} \sim 0.4$ MeV.

A number of specialized studies have been performed investigating the detection of Cherenkov light in scintillation detector mediums, including water-based liquid scintillator [280,281], slow-fluor scintillator [282], and linear alkylbenzene [283]. Two full-scale experiments have been able to perform a statistical measurement of Cherenkov radiation in a scintillating medium: Borexino, observing elastic scattering of sub-MeV solar ν_e in a large liquid scintillator detector [284], and ICARUS, observing cosmic muons in a large liquid argon detector [285]. However, no large-scale scintillation experiment has been able to observe Cherenkov light on an event-by-event basis. This is because the Cherenkov light output is typically smaller than the scintillation light output by around two orders of magnitude, which can be understood by comparing the numbers in figure 7-7a to the $\mathcal{O}(10^4)$ scintillation photons per MeV produced in liquid argon [158].

Event-by-event detection of Cherenkov light might be possible in CCM for a few reasons. First, liquid argon is transparent to UV light, which substantially increases the number of detectable Cherenkov photons as shown in figure 7-7a. TPB-coated PMTs can detect these UV Cherenkov photons by shifting them into the visible range. The wavelength-shifting efficiency of TPB actually increases by a factor of around two for $\lambda \sim 200$ nm compared to 128 nm argon scintillation light [286], potentially boosting the response of CCM PMTs to Cherenkov light. Second, the singlet lifetime of argon excimers is approximately 8 ns [160], which is much slower than the Cherenkov light emission timescale (< 1 ns) [287]. The 2 ns sampling time of the CCM digitizers may allow for the separation in time of prompt Cherenkov photons from the first scintillation photons. Finally, 20% of the PMTs in CCM are not coated in TPB. This means they are only sensitive to scintillation photons after they have been emitted from TPB somewhere else in the detector, delaying their arrival time in these PMTs. In contrast, Cherenkov photons in the visible regime will indeed produce signals on the uncoated PMTs. Though the larger UV component of the Cherenkov spectrum will not be accessible, even just a few photoelectrons detected by uncoated PMTs in the prompt time region would be a strong indication of Cherenkov light.

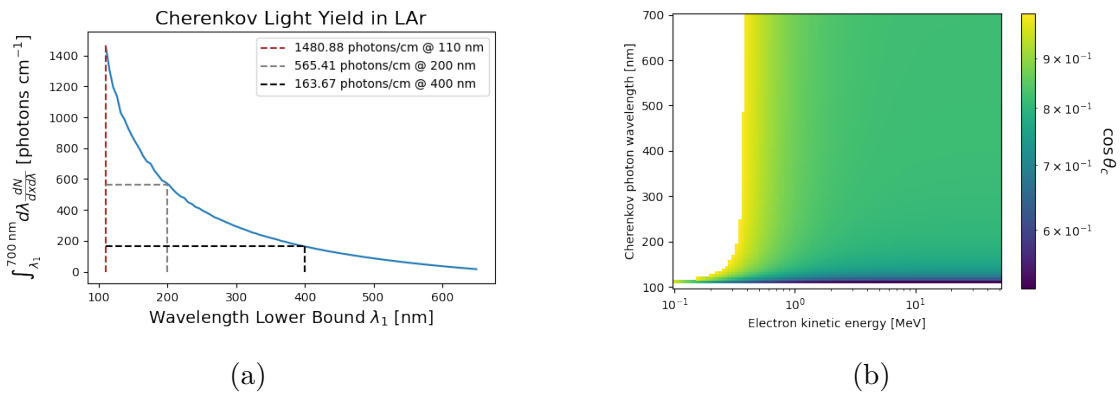


Figure 7-7: Figure 7-7a shows the integration in equation (7.1) as a function of λ_1 for $\lambda_2 = 700$ nm and $z = 1$. Figure 7-7b shows the Cherenkov cone angle $\cos \theta_C$ for an electron as a function of the photon wavelength and the electron kinetic energy.

7.2.1 Simulation-Based Sensitivity Estimation

In order to assess the sensitivity of CCM to Cherenkov light, we use the `Geant4` simulation of CCM200 described in Ref. [275]. We inject electrons at the center of the detector, traveling along the beam direction with kinetic energies of $\{1, 2, 3, 4, 5\}$ MeV. In addition to the usual scintillation light, Cherenkov light is emitted along the electron path within the 100 nm to 700 nm wavelength range. Both scintillation and Cherenkov photons are subject to wavelength-dependent effects within the detector, including absorption by impurities [164, 165] and Rayleigh scattering [167]. The TPB response in the simulation has not been carefully evaluated in the $200 \leq \lambda \text{ [nm]} \leq 400$ regime, so we have turned the TPB off for the purposes of this study. Instead, we apply a time offset for coated PMTs reflecting the TPB emission, which is sampled from an exponential with a time constant of 1.7 ns [288]. For both coated and uncoated PMTs, we apply an additional time offset related to the rise time of the PMT τ_{PMT} , which is about 2.5 ns in CCM. This is sampled from a Gaussian with a mean and width of $\tau_{\text{PMT}}/2$. For each photon that hits a PMT surface, we apply a weight reflecting the detection probability P_{det} , which is calculated as

$$\begin{aligned} P_{\text{det}}^{\text{uncoated}} &= P_{\text{PMT}} \\ P_{\text{det}}^{\text{coated}} &= P_{\text{PMT}} P_{\text{TPB}}, \end{aligned} \tag{7.2}$$

where P_{PMT} reflects the PMT quantum efficiency and P_{TPB} reflects the impact of TPB absorption and re-emission on the detection probability. The former is given by $P_{\text{PMT}} = 0.15$ for coated PMTs and

$$P_{\text{PMT}} = \begin{cases} 0.15 & 300 < \lambda \text{ [nm]} < 650 \\ 0 & \text{otherwise} \end{cases} \tag{7.3}$$

for uncoated PMTs, where 0.15 is the typical quantum efficiency of a cryogenic PMT [275, 289]. This reflects the approximation that TPB makes the coated PMTs sensitive to all photon wavelengths within 100 nm to 700 nm, while the uncoated PMTs are only sensitive to the visible range [289]. The TPB transmission probability

is approximated by

$$\begin{aligned}
P_{\text{TPB}} &= P_{\text{abs}}P_{\text{geo}}P_{\text{vis}}, \\
P_{\text{abs}} &= \begin{cases} 0.4 & \lambda < 200 \text{ nm} \\ 0.66 & \lambda > 200 \text{ nm} \end{cases}, \\
P_{\text{geo}} &= 0.4, \\
P_{\text{vis}} &= \begin{cases} 1.0 & \lambda < 400 \text{ nm} \\ 0.8 & \lambda > 400 \text{ nm} \end{cases},
\end{aligned} \tag{7.4}$$

where P_{abs} is a simple step function representation of the TPB absorption efficiency [286], P_{geo} is a geometric efficiency reflecting the fact that TPB emits photons isotropically off of the hemispherical PMT surface, and P_{vis} reflects the tendency of TPB to absorb visible light to some degree, as evidenced by its cloudy nature.

Using the setup described above, we simulate 10^4 electrons for each of the five tested kinetic energies to generate average templates of the expected number of registered hits in each PMT. In figure 7-8 we show templates for electrons with kinetic energies of $T_{e^-} = 1$ MeV and $T_{e^-} = 5$ MeV. Separate templates are shown including all photons and only scintillation photons (i.e., without Cherenkov light). The positions of each PMT here can be thought of as an unfolding of the cylindrical CCM detector, where the center of the image $((x, y) = (0, 0)$ in the middle panel) corresponds to the beam direction. Note that these templates only consider the first 8 ns after the electron starts traveling through the detector, as this is the region in which we expect Cherenkov light to stand out over scintillation light. This is indeed the case; the PMTs in figures 7-8a and 7-8c see much more light than those in figures 7-8b and 7-8d. This is especially true of the uncoated PMTs (dashed circles), which do not suffer the TPB detection inefficiency captured by equation (7.4) nor the time delay from TPB re-emission. The separation between the templates with and without Cherenkov light is more obvious for $T_{e^-} = 5$ MeV than for $T_{e^-} = 1$ MeV, indicating that Cherenkov light will be a stronger discriminator in the higher energy region. The

directional nature of Cherenkov light is also clear here. While scintillation photons are more-or-less isotropic, the event displays with Cherenkov photons included clearly show detected p.e. clustered around the beam direction. Figure 7-9 shows example event displays of the number of p.e. detected in each CCM PMT in a single simulated electron event for both $T_{e^-} = 1$ MeV and $T_{e^-} = 5$ MeV.

We use these simulations to build up a template-based Poisson likelihood reflecting the probability that a given event does or does not contain scintillation light. We evaluate a separate Poisson likelihood in each PMT across five time bins of 2 ns width each, which reflects the temporal sampling rate in CCM. The total likelihood is given by

$$\mathcal{L}_{\text{all/scint}} = \prod_{\text{PMTs } i} \prod_{\text{time bins } j} P_{\text{Poisson}}(k_{ij} | \mu_{ij}^{\text{all/scint}}), \quad (7.5)$$

where P_{Poisson} is the Poisson probability, k_{ij} is the observed number of p.e. in PMT i and time bin j , and μ_{ij}^{all} (μ_{ij}^{scint}) is the prediction for the templates with (without) Cherenkov light, as shown in figure 7-8. In the case that no events are predicted in a given PMT i and time bin j , we set $\mu_{ij} = 10^{-4}$, reflecting an uncertainty of one event over the 10^4 simulations. We can then use the detected p.e. in each simulation to build distributions of the log-likelihood-ratio test statistic

$$\Delta \log \mathcal{L} \equiv \log \mathcal{L}_{\text{all}} - \log \mathcal{L}_{\text{scint}}. \quad (7.6)$$


In figure 7-10 we show distributions of this test statistic for $T_{e^-} = 1$ MeV and $T_{e^-} = 5$ MeV, both with and without Cherenkov light. Given the overall sign of equation (7.6), $\Delta \log \mathcal{L}$ is larger when Cherenkov photons are included. This can be seen in figure 7-10, as the distributions with Cherenkov photons included sit at higher values of $\Delta \log \mathcal{L}$ compared to those without Cherenkov photons. The vertical line in this plot indicates the lower bound on $\Delta \log \mathcal{L}$ that can reject 99% of scintillation-only backgrounds, which can be thought of as, for example, beam-related neutrons depositing as many scintillation photons as an electron with the specified kinetic energy. Figure 7-11 shows curves of the efficiency to retain events with Cherenkov rings v.s. the rejection factor for events without Cherenkov rings, which are obtained by

considering successively larger lower bounds on $\Delta \log \mathcal{L}$. In the $T_{e^-} = 5$ MeV case, one can achieve a background rejection of 0.99 while retaining a signal efficiency of 0.99—thus, a Cherenkov likelihood cut has the potential to reduce backgrounds by around two orders of magnitude for free.


Of course, a few caveats must be mentioned regarding this study. Most importantly, we have only constructed a likelihood for electrons of a few example kinetic energies originating at the center of the detector and traveling along the beam direction. A full Cherenkov likelihood will need to be a function of the particle vertex (x, y, z) and direction (θ, ϕ) in addition to the particle energy. In order to properly assess the background rejection power of a Cherenkov-based cut, one will need to simulate neutrons in CCM using a realistic flux from the Lujan target after moderation by the CCM shielding. Finally, a more realistic TPB model within Geant4 itself will allow for a more robust prediction of the difference in coated and uncoated PMT response to Cherenkov light. That being said, the results from this simulation study indicate the power of a Cherenkov reconstruction algorithm in CCM and suggest that such a reconstruction is feasible for electrons with energies of a few MeV or greater. Angular cuts based on the reconstructed direction of the Cherenkov cone will help further reduce backgrounds by correlating events with the beam direction.

7.2.2 Identifying Cherenkov Light in Data


Having established the feasibility of a dedicated Cherenkov light reconstruction, we now discuss the path toward realizing this in CCM data. Cherenkov reconstruction requires a more detailed treatment of the timing information contained in PMT signals. The existing CCM reconstruction uses exponential smoothing, averaging, and derivative filters to determine regions of each PMT waveform with potential activity. A pulse is defined as a region in which the derivative goes below a given negative threshold, then becomes positive, then relaxes below the positive threshold, as shown in figure 7-12a. Figure 7-12b shows the identified pulses in an example CCM120 data event using this technique. Each pulse is then approximated by a triangle with a length and height given by the pulse length and absolute value of the integral of the




(a) With Cherenkov photons; $T_{e^-} = 1$ MeV



(b) Without Cherenkov photons; $T_{e^-} = 1$ MeV



(c) With Cherenkov photons; $T_{e^-} = 5$ MeV



(d) Without Cherenkov photons; $T_{e^-} = 5$ MeV

Figure 7-8: Templates of the average number of p.e. detected in each CCM PMT within the first 8 ns of an electron event. Different templates are shown for electron kinetic energies of $T_{e^-} = 1$ MeV and $T_{e^-} = 5$ MeV, both with and without Cherenkov photons. Coated (uncoated) PMTs are indicated by the solid (dashed) circles. Grey PMTs indicate those which registered no hits within the first 8 ns across all simulations. The dimensions on each axis are in units of cm.

waveform derivative over the pulse, respectively. Pulses were further required to be at least 20 ns in length to reduce the impact of noise. While this method was sufficient for the scintillation-based CCM120 analyses, which were mainly concerned with

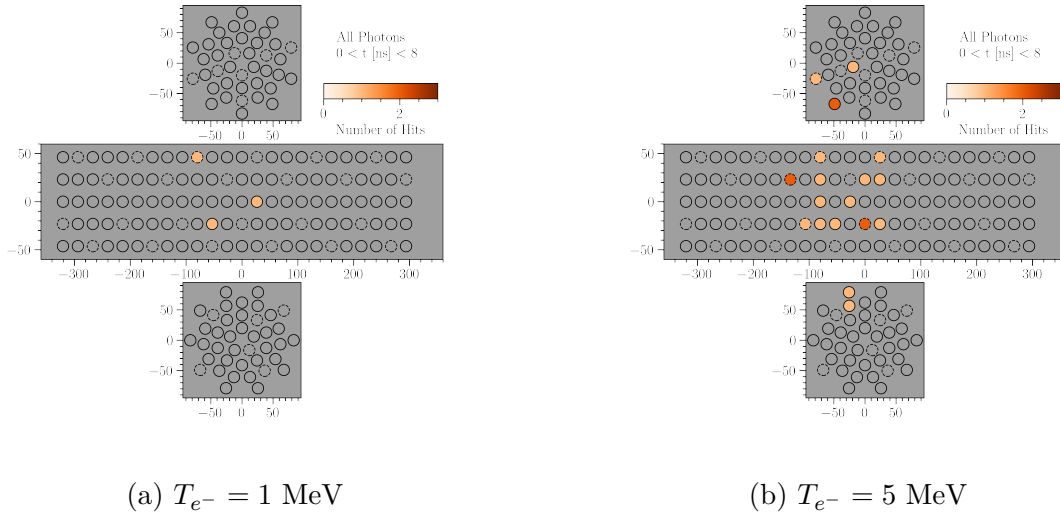


Figure 7-9: Example event displays showing the total number of p.e. detected in each CCM PMT within the first 8 ns of a single simulated electron event. Displays are shown for electron kinetic energies of $T_{e^-} = 1 \text{ MeV}$ and $T_{e^-} = 5 \text{ MeV}$. Coated (uncoated) PMTs are indicated by the solid (dashed) circles. Grey PMTs indicate those which registered no hits in the first 8 ns of this specific simulation.

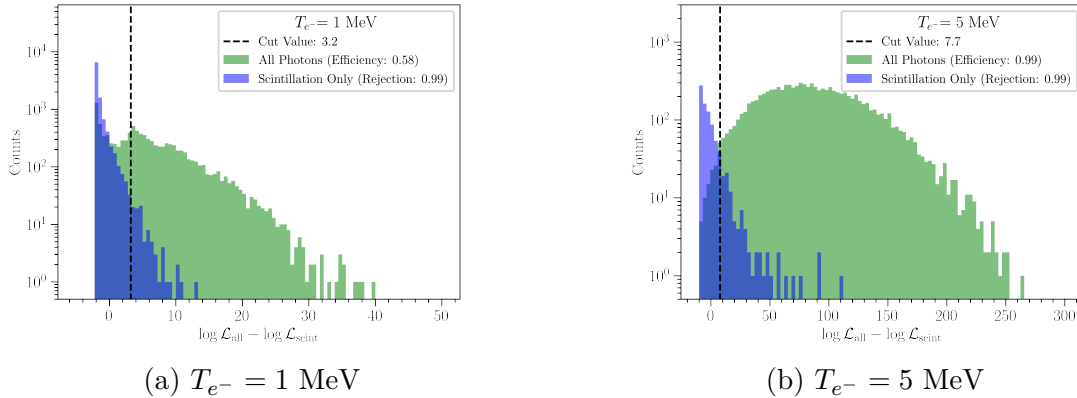


Figure 7-10: Distributions of the test statistic in equation (7.6) over all 10^4 simulations, considering either all photons or scintillation photons only. Distributions are shown for $T_{e^-} = 1 \text{ MeV}$ and $T_{e^-} = 5 \text{ MeV}$. The vertical line indicates the lower bound requirement which can reject 99% of scintillation-only backgrounds.

the many-p.e. scintillation pulse summed over all PMTs, identification of Cherenkov light will require a new pulse-finding algorithm that treats single p.e. pulses more carefully. It must also be able to retain pulses shorter than 20 ns to be sensitive to the Cherenkov-dominated early time window studied in section 7.2.1.




Figure 7-11: Curves of the efficiency to retain events with Cherenkov light (“Ring efficiency”) v.s. the fraction of events without Cherenkov light that can be rejected (“No-Ring Rejection Factor”), generated by considering successively larger lower bounds on $\Delta \log \mathcal{L}$. Different curves are shown for $T_{e^-} \in \{1, 2, 3, 4, 5\}$ MeV.

To address these issues, we have begun revamping the pulse-finding algorithm in preparation for CCM200. This includes the use of a more realistic single p.e. template to replace the triangle approximation. We use a parameterization of the single p.e. waveform based on that used by the IceCube collaboration [290],

$$w(t) = \frac{c}{[e^{-(t-t_0)/b_1} + e^{(t-t_0)/b_2}]^8}, \quad (7.7)$$

where c , t_0 , b_1 and b_2 are free parameters to be fit using data. To determine the best-fit single p.e. template for each PMT, we begin by isolating candidate pulses using the derivative filter in figure 7-12a. We restrict to candidate pulses with an amplitude between 20 and 40 ADC counts, which corresponds roughly to the single p.e. amplitude range in CCM. We fix two of the template parameters using the observed peak amplitude w_{\max} and peak time t_{\max} for each candidate pulse,

$$w_{\max} = \frac{c}{\left[\left(\frac{b_1}{b_2} \right)^{\frac{b_2}{b_1+b_2}} + \left(\frac{b_2}{b_1} \right)^{\frac{b_1}{b_1+b_2}} \right]^8}, \quad (7.8)$$

$$t_{\max} = t_0 + \frac{b_1 b_2}{b_1 + b_2} \ln(b_2/b_1).$$

Figure 7-13 shows the result of this template fit for candidate pulses on two different PMTs in CCM200 beam data. The next step is to determine a robust single p.e. template for each PMT in CCM200. These templates can then be used to unfold the observed waveform on every PMT for a given event into a series of single p.e. pulses at different times and with different amplitudes. Such a reconstruction will help pick out an excess of single p.e. pulses in the early time window, which would be an indication of Cherenkov light.

In order to inform the development of our Cherenkov reconstruction, we have isolated a new data sample in CCM200: cosmic muons. These are useful for two reasons: (1) downward-going cosmic muons will produce Cherenkov light along a known direction, and (2) isotropic emission of Michel electrons from muon decay-at-rest will help test the angular resolution of the Cherenkov reconstruction. Michel electrons can also serve as an important calibration point for the scintillation-based energy reconstruction, as they are emitted with a known energy distribution peaking at $m_\mu/2$ [172]. To tag cosmic muons, we have assembled six pairs of CosmicWatch detectors [291] on top of CCM. Each detector consists of a piece of plastic scintillator and a silicon photomultiplier, which detects photons produced in the scintillator. Two detectors can be stacked on top of one another and set to trigger on the coincident observation of light between the pair, indicative of a crossing cosmic muon. Figure 7-14a shows an image of the six CosmicWatch pairs on top of CCM. The detectors in each pair are separated by approximately 6 in. of foam to ensure that crossing cosmic muons are sufficiently downward-going to pass through the bottom of the detector. We have set up a dedicated cosmic muon trigger in CCM200 that saves an event whenever any of the six CosmicWatch pairs on top of CCM register a coincident signal. A diagram of this trigger is shown in section 7.2.2.

The cosmic muon trigger became operational in Summer 2022 and can run concurrently with the beam, strobe, and LED triggers. The total trigger rate is a bit under 1 Hz, meaning that $\mathcal{O}(10^6)$ cosmic muons have already been collected over the first few months of CCM200 data. Figure 7-15a shows the summed waveform across all PMTs in CCM200 for a single cosmic muon trigger. One can also see a delayed

signal from what is likely a Michel electron. Though we are still in the preliminary phase of analyzing data from the cosmic muon trigger, we touch briefly on a promising initial study of this dataset. This analysis examines the summed waveform in coated and uncoated PMTs separately in the top half and bottom half of the barrel of the detector. We look specifically at the initial rise time t_0 of each summed waveform, defined as the time tick at which the waveform passes 0.5% of its maximum value. In the scintillation only case, the coated PMTs should see a signal earlier than the uncoated PMTs ($t_0^{\text{coated}} < t_0^{\text{uncoated}}$), as only photons reflected from elsewhere in the detector will register a p.e. on the uncoated PMTs. This difference should diminish when Cherenkov light is involved, as both coated and uncoated PMTs will observe Cherenkov light to some level. We set the threshold to 0.5% as this reflects roughly the total number of Cherenkov photons generated by a crossing cosmic muon compared to scintillation photons. Section 7.2.2 shows the distribution of $t_0^{\text{coated}} - t_0^{\text{uncoated}}$ for PMTs in the top and bottom of the barrel over ~ 75 cosmic muon events. Given the angle of the Cherenkov cone emitted by a cosmic muon traveling downward along the central axis of the detector, only PMTs in the bottom half of the barrel will observe Cherenkov light. This is consistent with the distributions in section 7.2.2—the PMTs in the upper half of the barrel peak at $t_0^{\text{coated}} - t_0^{\text{uncoated}} \sim -5$ ns, while PMTs in the lower half of the barrel peak around zero. Thus, the delay of the signal in uncoated PMTs appears to disappear for PMTs sensitive to Cherenkov light. This is far from a definitive detection of Cherenkov light in the CCM detector, as one would need to carefully consider geometric effects in the CCM detector to predict $t_0^{\text{coated}} - t_0^{\text{uncoated}}$ distributions with and without Cherenkov light. However, it is a promising initial indication that Cherenkov reconstruction might be possible in CCM by leveraging differences in the coated and uncoated PMT signals.

7.3 Neutrissimos in CCM

We close our CCM discussion with an investigation of the sensitivity of CCM200 to the neutrissimo model of chapter 6. The idea of looking for dipole-coupled HNLs at

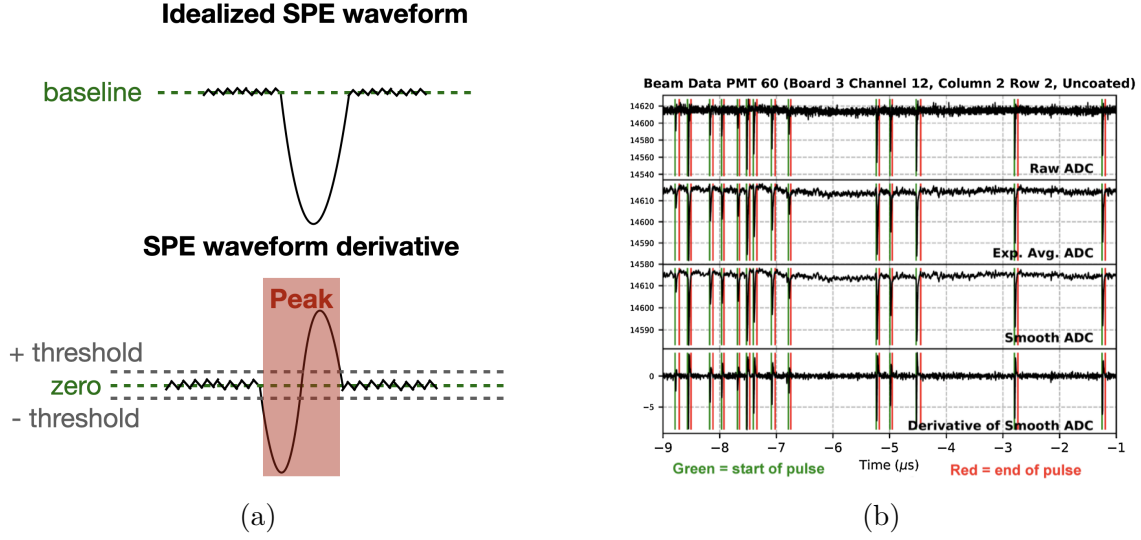


Figure 7-12: Figure 7-12a shows a schematic depiction of the derivative-based pulse definition in the current CCM reconstruction. Figure 7-12b shows an example of this pulse finder in a data event, from Ref. [30].





Figure 7-13: Two example waveforms from CCM200 beam data. The regions identified by the derivative filter are indicated in red and the result of the fit to equation (7.7) is indicated by the purple curves.

CE ν NS experiments has already been proposed in the literature [236]. The low energy threshold of these experiments enables a powerful signal definition: the coincidence of a nuclear recoil from upscattering ($\nu A \rightarrow \mathcal{N}A$) followed by a higher energy photon from neutrino decay ($\mathcal{N} \rightarrow \nu\gamma$). In particular, the upcoming NUCLEUS experiment is projected to have strong sensitivity to $d_{e\mathcal{N}}$ for $m_{\mathcal{N}} \lesssim 10$ MeV [236]. In this section, we investigate whether CCM200 might also have sensitivity to neutrinos in the $m_{\mathcal{N}} = \mathcal{O}(10$ MeV) regime. The important advantage of CCM200 is that it is



(a)

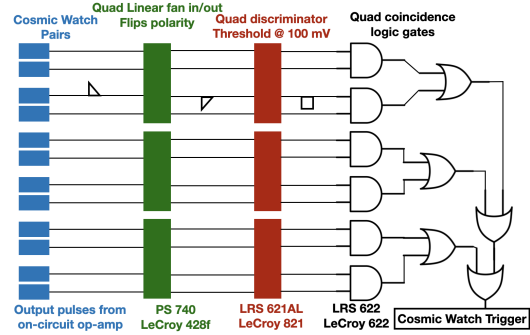



Figure 7-14: Figure 7-14a shows an image of the six CosmicWatch pairs on top of the CCM detector. Section 7.2.2 shows a schematic diagram of the cosmic muon trigger in CCM200.



(a)




Figure 7-15: Figure 7-15a shows the summed waveform across all PMTs in CCM for a single example cosmic muon trigger. The delayed signal from a Michel electron can also be seen. Section 7.2.2 shows the difference in rise times between coated and uncoated PMT signals in the top and bottom halves of the barrel of the detector (labeled “sides-top” and “sides-bottom”, respectively), as described in the text.

already running; projected sensitivities derived here will become limits within a few years.

The beam-dump source at Lujan is an ideal location to look for neutrissimos. The prompt ν_μ flux at CCM’s location is $4.74 \times 10^5 \nu/\text{cm}^2/\text{s}$ [30], corresponding to roughly $3.76 \times 10^{-2} \nu/\text{POT}$. These neutrinos can undergo Primakoff upscattering in the shielding around the target and along the flight path to the detector, creating a flux of neutrissimos. These neutrissimos can decay within the CCM detector to pho-

tons with $E_\gamma \sim 15$ MeV, which is relatively high energy compared to the background energy distribution [271]. A schematic depiction of this process is shown in figure 7-16. In this section, we define the neutrissimo signal as a single high energy photon in the detector. We do not evaluate our sensitivity to the coincidence channel discussed in Ref. [236], as the energy threshold of CCM does not yet enable CE ν NS detection. Additionally, though we only consider neutrissimo production via Primakoff upscattering, neutrissimos can also be produced by the $\pi^0 \rightarrow \gamma(\gamma^* \rightarrow \nu\mathcal{N})$ and $\gamma\nu \rightarrow \mathcal{N}$ channels within the target. Careful consideration of these additional detection and production channels may boost CCM’s sensitivity to neutrissimos, especially once the energy threshold is sufficiently lowered to enable a CE ν NS search.

As we rely on monoenergetic ν_μ with $E_\nu \sim 30$ MeV to undergo upscattering, we are sensitive to neutrissimos with $m_{\mathcal{N}} \lesssim 30$ MeV. We simulate neutrissimo upscattering and decay using the updated `LeptonInjector` [262] simulation framework. The public version of `LeptonInjector` includes a geometry file with a realistic description of the CCM200 detector and the surrounding Lujan facility, including the shielding shown in figure 7-16 and the TMRS shown in figure 7-1 as well as the concrete floor. Monoenergetic muon neutrinos are injected as an isotropic point source originating from the center of the lower tungsten target. We sample an upscattering location along the flight path to CCM, considering only a cone surrounding the CCM detector to improve the simulation efficiency. We then sample a neutrissimo decay location. If the neutrissimo path crosses the CCM200 fiducial volume, its decay is required to occur within this volume to further improve the simulation efficiency. We save the four-momentum and physical event weight (in units of POT $^{-1}$) for the final state photon. By scaling to the expected CCM200 exposure of 2.25×10^{22} POT, we can estimate the single photon event rate from neutrissimo decay in the total three-year CCM200 dataset. This `LeptonInjector`-based CCM200 neutrissimo simulation was presented to the public during the 2023 CCM Workshop [261].

To assess sensitivity, we consider the background estimation from the ALP search [271], as this analysis used a similar electromagnetic final state signal definition. A sizable reduction in the background rate is expected for CCM200 compared to CCM120; the

improved shielding already gives about an order-of-magnitude improvement [275], and the Cherenkov reconstruction simulation study described in section 7.2.1 suggests an additional two orders of magnitude in background reduction for a 5 MeV electron. In reporting results, we show a projected sensitivity range based on a background reduction within $[10^{-4}, 10^{-2}]$. The exact background reduction achieved by CCM200 will depend on the performance of the final Cherenkov reconstruction and argon filtration system as well as any improvements to the existing analysis cuts described in Ref. [271]. We assume a neutrissimo detection efficiency of 1.0, which is consistent with the expected ALP detection efficiency after CCM200 upgrades [275]. The photon energies are smeared according to an energy resolution of 15%, such that the prediction in each reconstructed energy bin μ_i is given by

$$\mu_i = \sum_{j=0}^{N_{\text{sim}}} \int_{E_i^{\text{low}}}^{E_i^{\text{high}}} \frac{w_j}{\sqrt{2\pi(0.15E_j)^2}} e^{\frac{-(E-E_j)^2}{2(0.15E_j)^2}} dE, \quad (7.9)$$

where N_{sim} is the number of simulated photons, E_i^{low} and E_i^{high} are the boundaries of reconstructed energy bin i , and w_j and E_j are the weight and energy of the j^{th} simulated photon from LeptonInjector. This energy resolution is relatively consistent with the CCM120 energy resolution at $E_\gamma \sim 20$ MeV [275] and a safe assumption once the updated Cherenkov reconstruction is in place.

We have run a series of simulations for neutrissimos with masses $m_{\mathcal{N}}[\text{MeV}] \in [1, 28]$ and dipole couplings $d_{\mu\mathcal{N}}[\text{GeV}^{-1}] \in [10^{-7}, 10^{-5}]$. Figure 7-17 shows the expected energy distribution of neutrissimo decay signal events and background events in CCM200 for $m_{\mathcal{N}} = 20.35$ MeV and $d_{\mu\mathcal{N}} = 3 \times 10^{-7}$ GeV $^{-1}$, considering a background reduction factor of 10^{-3} . The single photons from neutrissimo decay peak at $E_\gamma \sim 20$ MeV, as expected. We calculate the CCM200 sensitivity across parameter space using the χ^2 test statistic,

$$\chi^2(m_{\mathcal{N}}, d_{\mu\mathcal{N}}) = \sum_i \left(\frac{\mu_i(m_{\mathcal{N}}, d_{\mu\mathcal{N}})}{\sigma_i} \right)^2, \quad (7.10)$$

where $\mu_i(m_{\mathcal{N}}, d_{\mu\mathcal{N}})$ is the neutrissimo prediction in the i^{th} reconstructed energy in

and σ_i is the statistical \oplus systematic error on the background, appropriately scaled from Table 6.1 of Ref. [275]. We assume a χ^2 distribution with two degrees of freedom to draw projected sensitivities at the 95% confidence level (corresponding to a critical value of $\chi^2 = 6.18$). Figure 7-18a shows the preliminary projected CCM200 sensitivity to neutrino parameter space over the full three-year run. If we achieve a background reduction factor closer to 10^{-4} , CCM200 should be able to set world-leading constraints on the $d_{\mu\mathcal{N}}$ dipole coupling for $m_{\mathcal{N}} \sim 20$ MeV. As discussed in section 6.4, flavor-conserving UV completions of the dipole operator predict larger branching ratios for the decay $\mathcal{N} \rightarrow \nu_\tau \gamma$. As the neutrino single photon rate in CCM200 is primarily limited by the neutrino decay probability within the detector, CCM200 will be more sensitive to a scenario with such flavor-dependent dipole couplings. In figure 7-18b we show a preliminary estimate of CCM200’s sensitivity to the $d_{\tau\mathcal{N}} = d_{\mu\mathcal{N}} m_\tau / m_\mu$ discussed in section 6.4. The full CCM200 dataset will exclude a large region of the available parameter space in this flavor-dependent dipole model.

A few points are worth noting regarding these preliminary sensitivities. First, while the `LeptonInjector` simulation of CCM200 provides a relatively precise initial estimate of the neutrino decay rate in the detector, the geometric description of the Lujan facility will need to be verified before these sensitivities are ready for publication. This is especially true considering the strong dependence of the neutrino event rate on the exact shielding configuration between the tungsten target and the detector. That being said, any updates to the `LeptonInjector` CCM200 geometry made during this verification process are likely to be small and thus will not drastically shift the sensitivities in figure 7-18. The sensitivity might also improve after considering the additional neutrino production and detection mechanisms discussed earlier: $\pi^0 \rightarrow \gamma \nu \mathcal{N}$ and $\gamma \nu \rightarrow \mathcal{N}$ within the target, and nuclear recoil from Primakoff upscattering within the detector. These are all possible to simulate within `LeptonInjector`. The π^0 and γ flux were calculated for the ALP analysis [271], and nuclear recoil kinematics are already simulated within `LeptonInjector`. These channels will be carefully evaluated as the CCM200 neutrino analysis is prepared for publication. Further improvements to CCM200’s sensitivity can be made for modified

shielding configurations that introduce some amount of high- Z material, such as lead ($Z = 82$), in order to benefit from the Z^2 coherent enhancement of the Primakoff upscattering cross section [31]. Potential improved shielding configurations can and will be studied within `LeptonInjector`.

One must also consider the timing distribution of the photons from neutrissimo decay within the detector. CCM relies on strict timing cuts to remove neutron backgrounds; thus, it is important to ensure that the added travel time of non-relativistic neutrissimos does not push these photons outside of the CCM ROI. This is easily studied within `LeptonInjector`, as the entire path of the neutrino and neutrissimo is computed for each simulated event. Figure 7-19 shows the added time delay from the neutrissimo flight path for $m_{\mathcal{N}} = 20.35$ MeV (note that the time delay does not depend on the dipole coupling to first order). This is the most important neutrissimo mass to check, as it corresponds to the strongest projected constraints in figure 7-18. Two distinct populations can be seen at 5 ns and 26 ns, corresponding to upscattering in the shielding surrounding CCM and the TMRS, respectively. This suggests characteristic time delays $\Delta t \lesssim 30$ ns, which is comfortably within the 150 ns ROI of the ALP analysis [271]. Thus, the CCM timing cuts should not be an issue for the neutrissimo model.

The analysis presented here can in principle be used to set constraints using CCM120 data. However, the signal efficiency of neutrissimo events must be more carefully evaluated for this dataset. The CCM120 ALP analysis used a variety of cuts on the event length, reconstructed radius, and PMT charge uniformity to remove backgrounds. While the improvements made for CCM200 are likely to push the ALP (and thus neutrissimo) selection efficiency to greater than 0.95, the CCM120 ALP selection efficiency fell roughly within [0.2, 0.4] depending on the ALP energy [271]. Further, the CCM120 analysis used only the first 38 ns of each event to calculate the reconstructed energy [271], which will have a nontrivial impact on the assumed 15% symmetric energy resolution. Finally, the shielding configuration was updated between CCM120 and CCM200, and the CCM120 version is not currently modeled in `LeptonInjector`. We will need to study the impact of all of these effects on single

photon events before calculating CCM120 limits on the neutrissimo model.

The analysis presented in this section highlights the strong capability of CCM in constraining BSM physics. The high-intensity beam dump at the Lujan center can produce a significant number of potential BSM particles, while the extremely low duty factor helps isolate these particles from SM backgrounds. The large size, good energy resolution, and fast timing response of the CCM detector are well-suited for measuring any potential interactions of these BSM particles. The neutrissimo model discussed here is just one of many BSM scenarios on which CCM200 will place world-leading constraints [30, 145, 270, 271]. Additionally, the neutrissimo model specifically emphasizes the significant impact that Cherenkov light reconstruction will have on CCM200 sensitivity to BSM interactions with electromagnetic final states (which also occur in the ALP model).




Figure 7-16: Schematic depiction of prompt ν_μ from π^+ decay in the Lujan target upscattering to neutrissimos within shielding along the path to CCM200 and decaying to photons in the detector. The pink circle represents the TMRS shown in figure 7-1.




Figure 7-17: Figure 7-17a shows the distribution of background and signal prediction in CCM200 for $m_{\mathcal{N}} = 20.35$ MeV and $d_{\mu\mathcal{N}} = 3 \times 10^{-7}$ GeV $^{-1}$, considering a background reduction factor of 10^{-3} compared to CCM120. Section 7.3 shows the background-subtracted plot, with a red band indicating the expected statistical uncertainty on the background.




Figure 7-18: Figure 7-18a shows the expected sensitivity of CCM200 to the neutrino model, where the blue band corresponds to a background reduction factor between 10^{-4} (“CCM200 Low Background”) and 10^{-2} (“CCM200 High Background”). The MiniBooNE E_{ν}^{QE} allowed region (pink) and existing constraints (grey) come from Ref [31]. Figure 7-18b shows the same plot, but considering $\mathcal{N} \rightarrow \nu_{\tau}\gamma$ decays with $d_{\tau\mathcal{N}} = d_{\mu\mathcal{N}}m_{\tau}/m_{\mu}$.




Figure 7-19: The time delay of neutrissimo single photon decays with the CCM detector for the indicated mass and coupling, as calculated using `LeptonInjector`.

Chapter 8

Conclusions and Future Prospects

This thesis presents an array of experimental and phenomenological investigations into the nature of the MiniBooNE low energy excess (LEE). The MiniBooNE LEE, described in detail in chapter 2, refers to the 4.8σ excess of electron-like neutrino interactions observed by the MiniBooNE detector along the Booster Neutrino Beamline (BNB) [10]. This is one of the most significant and longest-standing anomalies with respect to the Standard Model (SM) of particle physics. The excess has remained consistent over the full 17-year dataset, corresponding to 18.75×10^{20} (11.27×10^{20}) protons-on-target (POT) in neutrino (antineutrino) mode. Due to the nature of a Cherenkov detector, MiniBooNE could not distinguish between electrons, photons, or collimated e^+e^- pairs—thus, any of these can contribute to the MiniBooNE LEE. Though SM explanations of the excess have been investigated thoroughly, both internally within MiniBooNE [10] and by the external community [80, 81], no promising SM candidates have emerged. Thus, the MiniBooNE LEE may be an indication of beyond-the-SM (BSM) physics. The most common BSM interpretation of the MiniBooNE anomaly is the 3 + 1 model, which predicts $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations at short baseline through an eV-scale sterile neutrino. This is an attractive solution as it would also explain the anomalous excess of $\bar{\nu}_e$ interactions observed by the LSND experiment, which was the original physics motivation behind the MiniBooNE experiment.

The MicroBooNE experiment was conceived to follow up on the MiniBooNE LEE.

MicroBooNE uses a liquid argon time projection chamber (LArTPC) detector [16], described in detail in chapter 3, to observe the interactions of neutrinos along the BNB. The detailed images of neutrino interactions produced by the LArTPC enable the separation of events with electrons and photons in the final state, allowing a suite of analyses looking for different explanations of the MiniBooNE excess. My research within MicroBooNE focused on the “two-body CCQE analysis”: a search for ν_e charged-current quasi-elastic (CCQE) interactions consistent with two-body scattering kinematics [148]. This analysis is described in detail in chapter 4. It uses a variety of techniques [24, 172], including several deep-learning-based methods [21, 22], to reconstruct and isolate signal events from backgrounds. The electromagnetic shower reconstruction, which drives the reconstruction of the original neutrino energy, is validated using two data-driven standard candles: the π^0 invariant mass peak and the Michel electron energy spectrum cutoff [172]. This study is presented in section 4.4.3. The most powerful technique used in the signal selection leverages an ensemble of boosted decision trees (BDTs) trained on a collection of variables describing the topology and kinematics of candidate ν_e CCQE interactions. This BDT ensemble is specifically tailored to select events consistent with a clean two-body interaction, which helps reduce the impact of complicated nuclear effects on our signal prediction.

The two-body CCQE analysis observed 25 events passing all signal selection requirements in the first 6.67×10^{20} POT of BNB data [148]. The statistical results from this dataset are presented in chapter 5. The two-body CCQE analysis does not observe significant excess of ν_e CCQE events at the lowest energies consistent with the MiniBooNE LEE. Using a two-hypothesis test and the CL_s test statistic, we are able to rule out the median expectation from the MiniBooNE excess at the 2.4σ confidence level. A signal strength scaling test performed using the Feldman-Cousins procedure sets an upper bound of 0.38 on the ν_e fractional contribution to the MiniBooNE LEE at the 2σ confidence level. The results obtained by our analysis are consistent with the other two MicroBooNE ν_e analysis—no analysis observes a significant low-energy excess of charged-current ν_e interactions like that predicted by

the MiniBooNE LEE [116].

Two external follow-studies regarding the MicroBooNE results are presented. The first of these is a combined fit of MicroBooNE and MiniBooNE data to the $3 + 1$ model [102], presented in section 5.3.1. This study shows that despite the strong constraints from MicroBooNE on ν_e interpretations of the MiniBooNE LEE, closed contours in $3 + 1$ parameter space still exist at the 3σ confidence level in the combined MiniBooNE-MicroBooNE fit. Thus, the MicroBooNE results do not definitively rule out the $3 + 1$ interpretation of the MiniBooNE excess. The second study, presented in section 5.3.2, looked into the implications of a $\bar{\nu}_e$ explanation of the MiniBooNE excess [201]. Due to the isoscalar and non-isoscalar nature of the carbon and argon nucleus, respectively, the $\bar{\nu}_e$ cross section tends to be much more suppressed at low energy in argon compared to carbon, while the same is not true for ν_e interactions. Thus, MicroBooNE will be less sensitive to the MiniBooNE LEE if it comes from $\bar{\nu}_e$ interactions. Ref. [201] shows that the MicroBooNE data are consistent at the 2σ confidence level with a scenario in which the MiniBooNE excess is comprised entirely of $\bar{\nu}_e$ events.

In chapter 6 we discuss a mixed model comprising an eV-scale sterile neutrino and an MeV-scale heavy neutral lepton with a dipole coupling to active neutrinos, referred to as a neutrissimo \mathcal{N} [27, 31]. It is shown that this model can significantly relax the internal tension in global $3 + 1$ by explaining the bulk of the MiniBooNE excess through neutrissimo decays to single photons. As discussed in section 6.3, the neutrissimo with $m_{\mathcal{N}} \sim 500$ MeV can simultaneously provide a good explanation to the energy and angular distributions, which is not possible within the $3 + 1$ model alone. In Ref. [31], we have calculated world-leading constraints on the neutrissimo model using the high dE/dX sideband of the MINER ν A elastic scattering analyses. This study is presented in section 6.4. While the MINER ν A constraints are indeed very strong in the $m_{\mathcal{N}} = \mathcal{O}(100 \text{ MeV})$ region, they do not rule out the MiniBooNE-preferred region in neutrissimo parameter space at the 95% confidence level. This is because the kinematic cuts in the MINER ν A analysis remove a majority of neutrissimo events for $m_{\mathcal{N}} \sim 500$ MeV. A dedicated MINER ν A analysis would likely be

able to make a definitive statement on the neutrissimo explanation of the MiniBooNE LEE.

Finally, we close with the Coherent CAPTAIN-Mills (CCM) experiment in chapter 7. CCM uses a liquid argon detector at the Lujan proton beam dump facility of the Los Alamos Neutron Science Center (LANSCE) to look for the interactions of potential beyond-the-Standard-Model (BSM) particles. The CCM detector uses an array of photo-multiplier tubes (PMTs) to observe the scintillation light produced in these interactions. CCM leverages the low duty factor of the Lujan beam dump to make strict timing cuts on signals coming from the beam, isolating prompt signal events corresponding to BSM particles from slower beam-related neutron backgrounds. This strategy has allowed CCM to place strong limits on a wide variety of BSM scenarios, including light vector-portal dark matter [30], leptophobic dark matter [270], and axion-like-particles [271], using only data from a six-week 120-PMT engineering run. The full 200-PMT detector, “CCM200”, started taking data in Summer 2022 for its nominal three-year run at the Lujan facility. CCM200 will make a number of improvements upon the 120-PMT engineering run, including upgraded shielding and the deployment of an argon filtration system, which will help set world-leading limits on these models.

One of the most promising improvements is the development of a Cherenkov light reconstruction algorithm, which can help separate BSM events with electromagnetic final states from SM backgrounds. Section 7.2.1 presents a simulation-based study that indicates the feasibility of Cherenkov reconstruction in CCM thanks to the combination of PMTs coated and uncoated with wavelength-shifting tetraphenyl butadiene. The study suggests that uncoated PMTs will be especially sensitive to Cherenkov light in the visible regime for the first ~ 10 ns of a given electromagnetic event. Section 7.2.2 discusses some of the initial steps toward developing the Cherenkov reconstruction in data, including the deployment of a dedicated cosmic muon trigger and improvements to the CCM pulse finding algorithm. We discuss a potential use-case for the Cherenkov reconstruction in section 7.3, which presents a calculation of CCM200’s sensitivity to the neutrissimo model discussed above. Prelim-

inary estimates suggest that CCM200 will be sensitive to new regions of neutrino parameter space for $m_{\mathcal{N}} \sim 20$ MeV over its full three-year dataset.

Neutrinos have never failed to surprise us. Since the first detection of neutrinos in 1956, experimental anomalies like the solar neutrino problem and the atmospheric neutrino anomaly have provided a guiding light toward a more complete understanding of the “little neutral one”. The experimentally-established fact that neutrinos have nonzero mass is not predicted by the Standard Model, and the comparatively tiny neutrino mass with respect to other particles is a strong motivation for BSM physics. Anomalies in short baseline neutrino experiments could very well be hinting that the neutrino has secrets yet to reveal. This thesis takes a few concrete steps along the path lit by the MiniBooNE anomaly. Only time—and clever experimental analyses—will tell where this path leads.

Appendix A

Publication: *Convolutional neural networks for shower energy prediction in liquid argon time projection chambers*

My first research project as a graduate student involved the development of a convolutional neural network for electromagnetic shower energy reconstruction in MicroBooNE. While the MicroBooNE network never saw the light of day, the network was adapted for use in general LArTPC experiments by an MIT undergraduate student, Kiara Carloni. Kiara led the study discussed in the *JINST* publication below [23]. I served mainly as a mentor for this project, as it was extended from my early MicroBooNE work.

PAPER

Convolutional neural networks for shower energy prediction in liquid argon time projection chambers

To cite this article: K. Carloni *et al* 2022 *JINST* **17** P02022

View the [article online](#) for updates and enhancements.

You may also like

- [Neutrino–nucleus cross sections for oscillation experiments](#)
Teppei Katori and Marco Martini
- [Measurement of space charge effects in the MicroBooNE LArTPC using cosmic muons](#)
P. Abratenko, M. Alrashed, R. An et al.
- [Ionization electron signal processing in single phase LArTPCs. Part II. Data/simulation comparison and performance in MicroBooNE](#)
C. Adams, R. An, J. Anthony et al.

RECEIVED: November 4, 2021

REVISED: January 3, 2022

ACCEPTED: January 27, 2022

PUBLISHED: February 11, 2022

Convolutional neural networks for shower energy prediction in liquid argon time projection chambers

K. Carloni,* N.W. Kamp, A. Schneider and J.M. Conrad

*Department of Physics, Massachusetts Institute of Technology,
Cambridge, MA 02139, U.S.A.*

E-mail: kcarloni@mit.edu

ABSTRACT: When electrons with energies of $O(100)$ MeV pass through a liquid argon time projection chamber (LArTPC), they deposit energy in the form of electromagnetic showers. Methods to reconstruct the energy of these showers in LArTPCs often rely on the combination of a clustering algorithm and a linear calibration between the shower energy and charge contained in the cluster. This reconstruction process could be improved through the use of a convolutional neural network (CNN). Here we discuss the performance of various CNN-based models on simulated LArTPC images, and then compare the best performing models to a typical linear calibration algorithm. We show that the CNN method is able to address inefficiencies caused by unresponsive wires in LArTPCs and reconstruct a larger fraction of imperfect events to within 5% accuracy compared with the linear algorithm.

KEYWORDS: Pattern recognition, cluster finding, calibration and fitting methods; Noble liquid detectors (scintillation, ionization, double-phase); Time projection Chambers (TPC)

ARXIV EPRINT: [2110.10766](https://arxiv.org/abs/2110.10766)

*Corresponding author.

Contents

1	Introduction	1
2	Methods	2
2.1	Dataset	2
2.2	Clustering-based linear algorithm	4
2.3	Convolutional neural networks	4
3	Experiments and results	7
3.1	Size comparison	7
3.2	Loss function comparison	9
3.3	Network input comparison	10
3.4	Comparison to the linear algorithm’s performance	14
4	Conclusion	17
A	Architecture details	18
B	Further detail on the studies	18
C	Systematic variation study	20
C.1	Flat noise	20
C.2	Gaussian per-wire noise	22

1 Introduction

Liquid argon time projection chambers (LArTPCs) [1] consist of a large volume of liquid argon, a drift region within the argon where a static electric field is maintained, and planes of anode wires where charge is measured and collected. A charged particle traversing the detector will ionize the liquid argon, leaving a trail of ionization electrons along its path. The static electric field is used to drift these ionization electrons. The anode planes consist of a series of independent wires on which the arriving charge is measured. This segmented readout provides two dimensional charge information (per wire and per unit time) from each anode plane. Thus, the data coming from a given anode plane can be thought of as an image, in which each “pixel” corresponds the amount of charge measured on a single wire over a specified time interval. Figure 1 gives a schematic summarizing the operational principle behind a LArTPC.

LArTPCs are an increasingly common detector technology used to measure neutrino interactions [2–5]. At energies of $\mathcal{O}(100)\text{MeV}$ or greater, electrons and photons produced in such interactions will undergo a chain of bremsstrahlung radiation ($e^\pm \rightarrow e^\pm\gamma$) and pair production $\gamma \rightarrow e^+e^-$ known as an “electromagnetic shower”. Typical algorithms for the energy reconstruction

of such electromagnetic showers in LArTPCs have two main steps [6–9]. First, pixels are identified in the anode planes that measured charge associated with the shower of interest. For example, this can be done by finding a cone that contains the shower and computing the total charge within the cone, excluding pixels with a low probability of being related to the shower [10]. Second, the total charge is linearly mapped to the shower energy. These traditional clustering-based linear algorithms are fast, straightforward, and accurate for most LArTPC shower events.

However, when a LArTPC detector has sections of unresponsive wires, a subset of a shower’s charge will pass undetected and the linear algorithm will underpredict the shower’s energy. This effect is more likely in higher energy showers, since their charge is spread over a larger area and has a higher chance of passing through an region of unresponsive wires. In this case, the simple linear mapping to shower energy has difficulty correcting for the missing charge simultaneously across events of all energies, and events are often mis-reconstructed to lower energies.

A convolutional neural network (CNN) trained on a dataset with a substantial proportion of unresponsive wires could learn to reconstruct the energy of events with high fractions of missing charge. CNNs train kernels to recognize patterns within an image, and thus efficiently use their parameters to extract information from very large image inputs. Visual features such as the shower size, radial extent, and longitudinal structure contain information about the shower’s energy, so neural networks trained to recognize these patterns can potentially correct for the nonlinear mapping between charge and shower energy. Additionally, the position of unresponsive wires can be directly provided as input to the neural network to enable even more accurate reconstruction.

This study investigates the energy reconstruction performance of CNN models with varied sizes, architectures, training parameters, and inputs. We then compare the top CNN models’ performance to that of traditional energy reconstruction methods under different detector conditions. We show that the CNN method is able to address inefficiencies caused by unresponsive wires in LArTPCs and reconstruct a larger fraction of imperfect events.

2 Methods

2.1 Dataset

We used the PILArNet public simulation pipeline [11] to generate simulated LArTPC images. We assume a detector setup that includes three parallel anode planes with wires, spaced 3 mm apart, oriented at 60° (U), -60° (V), and 0° (Y) with respect to the zenith. The PILArNet simulation process first instantiates a collection of particles at a vertex with randomly distributed kinetic energies, then tracks their charge deposition through the detector volume using GEANT4 [12]. The 3D charge information is voxelized into 3mm length cubes and recorded, along with summary information on every particle generated, in a sparse data file. Each pixel in PILArNet images has an intensity value intended to represent the charge deposited in that region of the TPC. Hereafter we refer to this intensity value as the charge “Q”. The PILArNet dataset is described in full in ref. [13].

The PILArNet procedure produces 3D voxel data for monodirectional electron showers. However, a LArTPC reads out charge information projected onto the anode planes through the drift process, and we are interested in showers of all orientations. To adapt the PILArNet dataset to our needs, we rotate and project each shower event onto the anode planes to get a collection of isotropically




Figure 1. Diagram showing the operational principle behind a LArTPC. Charged particles traverse the argon volume, leaving a trail of ionization electrons which are drifted through an electric field to a series of wire planes. Figure adapted from ref. [3]

oriented LArTPC events as they would be readout by the U, V, and Y anode planes. Every event is assigned a random rotation, which is composed with a 60° , -60° , or 0° degree rotation in the X-Y plane to get the three anode plane views. For each rotation applied to an event, the charge in every input voxel is distributed to adjacent voxels in the rotated voxel grid with a weight inversely proportional to the euclidean distance between voxel centers. This method is simple and preserves the total charge in the image. Since the voxels are small, the particular choice of reapportionment metric is not significant. Finally, the two dimensional anode plane images are created by summing the charge from voxels along the axis perpendicular to the wire planes.

Our resulting simulated LArTPC images are comprised of 600 k single-electron shower events of energies uniformly distributed between 0 and 1000 MeV with a uniform angular distribution. Uniform distributions were chosen to minimize any energy based or angular bias of the predictors. We divided our dataset into 400 k training events and 200 k validation events. Each event in the samples contains three base 768×768 images, which are the U, V, and Y plane readouts. An example event image is displayed in figure 2.

We further augment our base dataset by adding unresponsive wires. First, a set of unresponsive wire chunks is placed randomly within the full extent of the anode planes with 20 k, 20 k, and 21.5 k wires composing each plane (these dimensions match the planned size of the DUNE far detector [14]). The unresponsive wire chunk lengths were randomly sampled to be consistent with typical ASIC sizes in a LArTPC: most chunks were less than 20 wires long, and all were at most 70. The total




Figure 2. An example electron shower event (dark purple) from our testing dataset. The locations of unresponsive wires are marked by superimposed pink stripes.

fraction of unresponsive wires was fixed to 10 % for each plane, which is the magnitude reported by a 2017 MicroBooNE analysis of noise features in their LArTPC [15]. We then assigned each event a location within the detector volume and zeroed out all charge corresponding to unresponsive wires.

2.2 Clustering-based linear algorithm

The typical clustering algorithm approach to energy reconstruction identifies the shower of interest, and linearly calibrates the charge associated with the shower to a shower energy. Our simulated dataset contains only single electron shower events, with no background. Thus in our implementation of the traditional energy reconstruction algorithm, we linearly map the total charge in the Y collection plane image to a shower energy, bypassing the clustering step.

To find our linear mapping parameters, we first separate our event sample into 50 bins in total-Y-plane-charge, each 20 Q wide. Within each total-charge bin the distribution of true shower energies is fit to a Gaussian distribution. Any bins for which the Gaussian fit is sufficiently imprecise are skipped: that is, the bins are skipped if the fit’s uncertainty on a parameter is greater than ten times the parameter’s value. The remaining bins are interpreted as observations of the true shower energy at specific total charges with corresponding errors. These data points are then used as input to a least squares fit of the linear mapping between the binned shower charge and true shower energy. On datasets with unresponsive wires, we found this procedure corrected for bias in all energy bins over 60 MeV, but did not do so in bins below.

2.3 Convolutional neural networks

Different neural network architectures can be better suited for different problems. For this study, we focused on two neural network architectures: the residual network [16] and the inception network [17]. Each has found success in previous LArTPC deep learning projects [18, 19]

The residual network (ResNet [16]), introduced in 2015, was designed to allow deep layers in the network to learn small adjustments to the final output easily. Its central structure is the residual block, a network layer whose output is a sum of its input and a filtered “residual.” The original authors suggested that because residual blocks make it easy for deep layers to implement an identity map, very deep ResNets can converge more quickly to more accurate solutions.

The first version of the inception network [17] was published in 2014 and designed as a wider, rather than deeper, network. For the purposes of this paper we have labeled this architecture InceptNet. Its fundamental structure is the inception block, which combines the results of convolutional filters of various sizes into one output. Inception blocks are therefore suited to describing features of different sizes simultaneously. The original authors suggested that such hybrid blocks could approximate a sparse network structure by dense component calculations. Thus, this architecture might work well for LArTPC images, which are sparse by nature. A variation of the inception network was successfully implemented for energy reconstruction by the NO ν A collaboration in 2019 [18].

The original authors of the inception network have since released three updated versions. These updates employ new structural efficiencies to express the same operations in fewer calculations, and thus can form even deeper networks. However, these architectures quickly grow too large for our purposes, so for this paper we focus solely on the original version.

We implemented both the ResNet and InceptNet architectures according to the same general structure as the original works. These are, respectively, depicted schematically in figures 3 and 4. Each network has an initial gate, which scales down the image spatial dimensions and increases the feature dimensions. The gate is then followed by a stack of convolutional layers, which constitute the bulk of the network. For each architecture, we built models of four different sizes, “Small,” “Medium,” “Large,” and “Huge,” by increasing the depth of this stack. Each network size has roughly an order of magnitude more trainable parameters than the previous. We scaled architectures up according to two main principles: first, by stacking on additional layers of the architecture’s basic block, and second, by adjusting the feature map sizes throughout the network so that the number of parameters in sequential layers increases gradually, the input image dimensions decrease gradually, and the feature dimension increases. The exact implementation of each network layer is most likely not significant [17]. Details of the four network sizes are given in appendix A, and figures 17 and 18. The quantitative properties of each model size, including the exact numbers of parameters, are listed in table 1 and are discussed later in this work. The output of the convolutional layers is collapsed into a vector by an average pool layer. This vector can then optionally be combined with additional vector inputs, and is finally passed through a dense fully-connected layer, which outputs the predicted energy.

The networks take as input a batch of layered images and vectors of additional information. The primary inputs were the U, V, Y output planes of the LArTPC detector, and a three-component vector containing the total recorded charge on each plane. We also investigated whether the networks could utilize information about unresponsive wire placement. In one set of tests, we fed in images with six channels: three base U, V, Y 768×768 pixel images, and three 768×768 images indicating the locations of dead wires for the corresponding planes.

CNNs are trained using stochastic gradient descent to iteratively adjust the network’s weights such that its performance is improved with respect to some loss function [20]. To reduce memory costs associated with large image-based datasets, during optimization the training set is sampled in smaller batches. The model’s weights are thus adjusted in a series of less-confident steps, rather than one sure leap. The weight adjustment step for each batch can be further modified to incorporate information about the previous adjustment: the parameters at batch $t + 1$ are related to those at batch t by $p_{t+1} = p_t - m \cdot u_t + \ell \cdot g_{t+1}$, where g_{t+1} is the gradient for batch $t + 1$, $u_t = p_{t-1} - p_t$ is the adjustment factor for batch t , and ℓ, m are fixed hyperparameters called the learning rate and

ResNet Architecture:

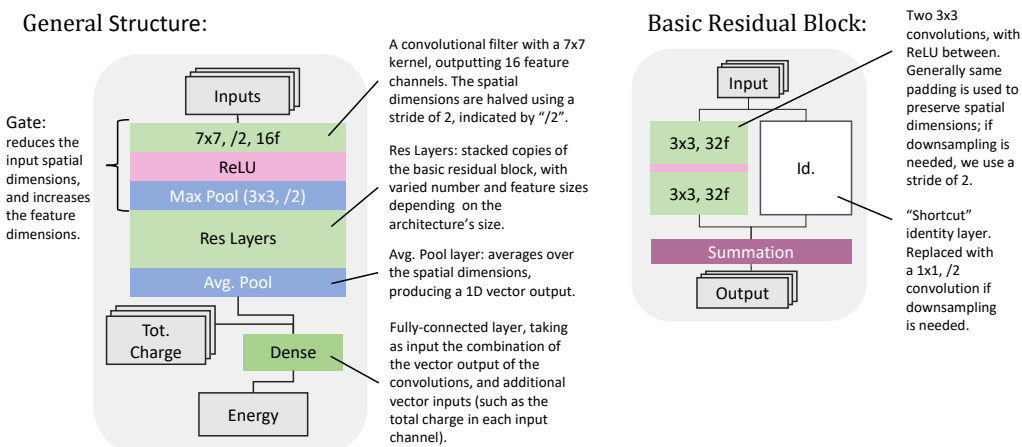


Figure 3. The basic structure of a ResNet consists of a starting gate, a stack of residual layers, and then a fully-connected layer which decodes vector output into a single energy prediction (*left*). Each residual layer is formed out of stacks of the basic residual block (*right*) which transforms its input by adding a small convolutional correction.

InceptionNet Architecture:

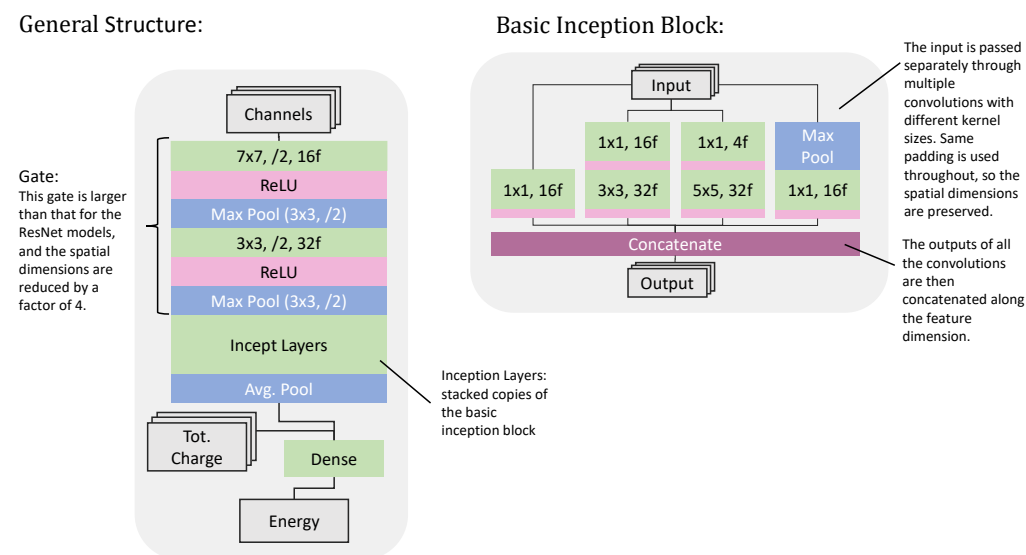


Figure 4. The basic structure of our InceptionNet is similar to that of the ResNet, but includes a slightly larger gate (*left*). Inception layers consist of stacks of the basic inception block, which concatenates the output of convolutional filters of many different sizes (*right*). Concatenation is possible since all convolutional filters use the same padding, and thus preserve spatial input dimensions.

momentum respectively. We tested different loss functions to explore how they would affect the final network performance at different energies.

We implemented and trained our networks using PyTorch’s software library [21], and we optimized our weights using PyTorch’s stochastic gradient descent algorithm with an initial learning rate of $\ell = 0.1$ and a momentum factor $m = 0.5$. Since we expected training to converge on some local optimal weight configuration, we reduced our learning rate when performance improvements plateaued: if for ten epochs the total loss did not decrease below 0.9999 times the best value, we reduced the learning rate by a factor of 10.

Training was divided into epochs. In each epoch, the network learned from a randomized subset of 160 k images from the 400 k event training dataset, and validated its results on a randomized subset of 6 k images from the 200 k event testing dataset. Due to time and GPU memory constraints, we trained all models for 100 epochs total. We then tested model performance using the trained parameters from the 100th epoch, which were generally the most refined. We expect all the models would see additional small performance gains if trained further, especially those using the Residual architecture.

3 Experiments and results

In this section, we compare the performances of the different networks described in the previous section. In section 3.1, we examine the performance as a function of network size. Next, in section 3.2, we compare networks trained using three different loss functions. In section 3.3, we evaluate the network performance with and without including information on unresponsive wires. Finally, in section 3.4, we compare the performance of the neural networks to that of a linear reconstruction algorithm. See appendix B for more details on the studies described in this section.

3.1 Size comparison

The modular designs of the ResNet and InceptNet architectures imply that we must choose the size of the network. In general, an increased number of training parameters will allow the network to capture more complex structure. However, larger network sizes come with a substantial increase in the training time and memory consumption, and so it is worth investigating the reconstruction performance as a function of network size. Since LArTPC images are relatively sparse compared to other datasets like photographic images, we expect that increasing the network size should have diminishing returns [18].

We tested models of four different sizes, “Small”, “Medium”, “Large”, and “Huge”, as described in 2.3. Table 1 contains the quantitative properties of each size model. The number of parameters scaled linearly with the space in memory required to store the model, and non-linearly with the training and evaluation time needed per image.

All models were trained on training and evaluation sample datasets of 400 k and 200 k images each, using the setup described above with a linear (L1) loss function. In this comparison, we used the datasets with unresponsive wires for both training and testing. In order to use the memory of our GPUs efficiently, we trained the huge model with a batch size of 64 images, and all the other sizes 128. The loss curve for every model, depicted in figure 5, shows that substantial learning occurs by epoch 40. After this point, the loss curves show slight improvements at about every ten epochs, when the learning rate decreased. We expect training for additional epochs could yield small additional performance improvements.

Table 1. Quantitative properties of different-sized models. Each increase in model size corresponded to an order of magnitude increase in the number of parameters.

Architecture	Size	Number of Parameters [M]	Training Time / Image [ms]	Validation Time / Image [ms]
InceptNet	Small	0.088	2.09 ± 0.05	0.60 ± 0.003
InceptNet	Medium	0.915	2.71 ± 0.49	0.73 ± 0.002
InceptNet	Large	5.798	4.53 ± 0.46	1.18 ± 0.002
InceptNet	Huge	21.361	8.88 ± 0.68	2.25 ± 0.004
ResNet	Small	0.082	2.50 ± 0.34	0.66 ± 0.002
ResNet	Medium	0.705	3.51 ± 0.38	0.85 ± 0.0001
ResNet	Large	4.378	4.61 ± 0.04	1.07 ± 0.0001
ResNet	Huge	23.786	6.12 ± 0.80	1.44 ± 0.003

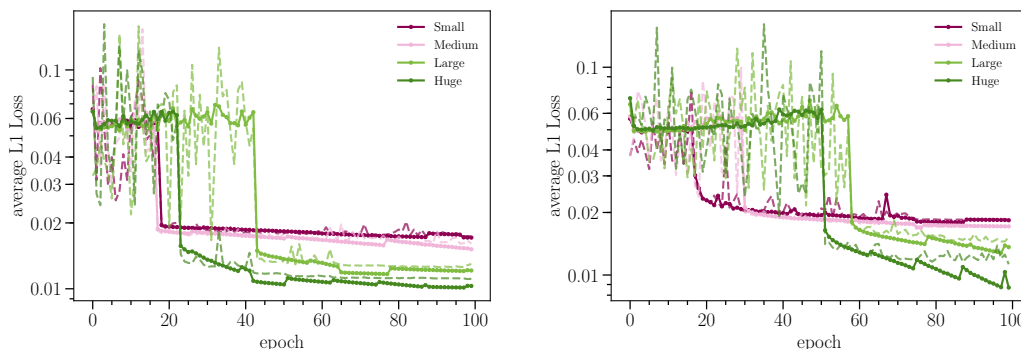


Figure 5. The average training (solid) and testing (dashed) L1 loss per epoch, for the different sized InceptNet (*left*) and ResNet (*right*) models. 10-epoch periodic dips in the training loss coincide with scheduled reductions of the learning rate.

As one would expect, the uncertainty on the reconstructed energy decreases as the network size increases. For the InceptNet Small, Medium, Large, and Huge models, the top 68 % of events have a reconstructed energy within 3.96 %, 3.86 %, 3.15 % and 2.79 % of the true value respectively, and for the corresponding ResNet models it is within 3.89 %, 3.48 %, 3.40 %, and 2.56 %.

Figure 6 depicts the performance of each model size as a function of the true shower energy. In the upper plots, we see that the modes of the fractional error distributions are stable and close to zero across all energies. Both the InceptNet and ResNet large models show an at least 1 % overprediction bias at all energies. In the lower plots, we see that the widths of the distributions improve across all energies with increasing model sizes. All models reconstruct the lowest energy bin, 0 – 100 MeV much less precisely than higher energy bins.

The Large networks show a visible overprediction bias compared to the other sized networks across all energy bins, and a slightly larger overall overprediction bias (0.7%, as opposed to 0.5% or less). We did expect over prediction to be a more likely outcome than under, since for a large subset

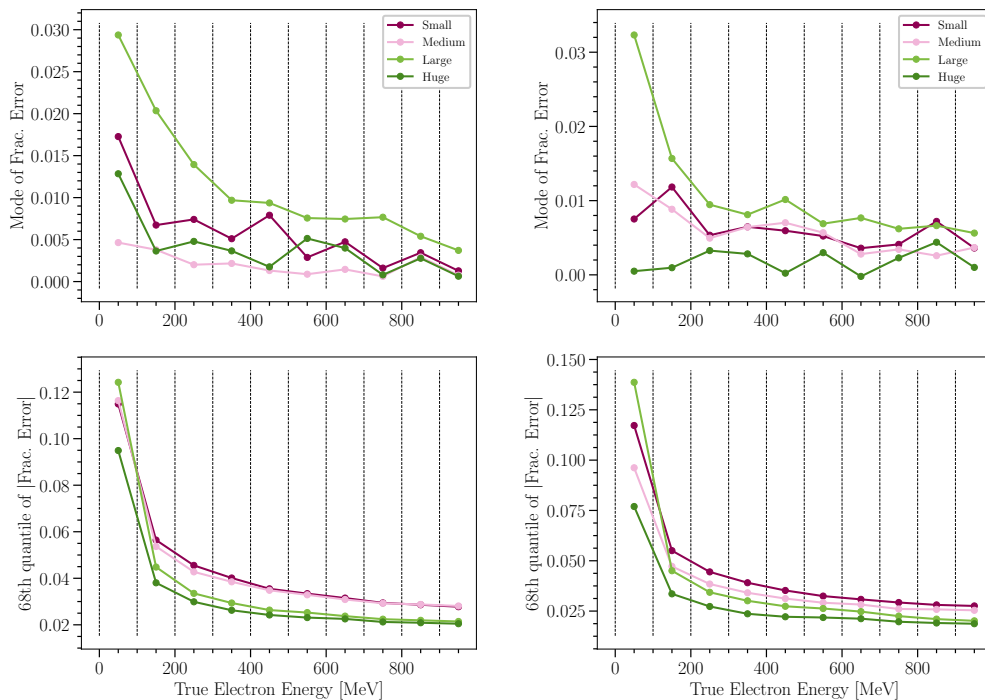


Figure 6. The mode and absolute 68th quantiles of the fractional reconstruction error, in each bucket of true shower energy, for the InceptNet models (*left*) and the ResNet models (*right*).

of events, the models need to compensate for missing charge during reconstruction. Ideally, this compensation would be regulated by factors such as large chunks of discontinuity in the shower or the known presence of unresponsive wires, but it’s likely that a flat correction is also used. If this were the case, then the error (mode and spread) would be worst in the lowest energy bin, which is what we see in figure 6. We think it is unlikely that this bias is a feature of solely the Large sized network. The size of this effect may be an accident of the final state parameters, and could be corrected after additional epochs of training.

The Huge models produced the narrowest, most sharply peaked distribution across the full range of energies. However, Huge models are computationally very expensive. Not only is the per-image training time 150 – 200 % that of the Large models, but also the quadrupled model size limits the maximum number of events that can be batched together, so the training procedure needs to cycle smaller sets of events on and off the GPU. For other applications, the performance improvement may be worth the extra cost in time and memory — for this study, the large models offer excellent predictive power, with small variance in fractional error, for a reasonable cost.

3.2 Loss function comparison

For the next study, we compared the performance of networks trained with a linear (L1), fractional (Frac.), and mean squared error (MSE) loss function. A linear loss function weights every event

equally, while a squared error loss function weights outlier reconstructions more heavily. The fractional loss function weights errors in low energy events more heavily, and using it to train models could prioritize the performance in those events. Fractional loss was also used successfully by the NOvA collaboration [18]. For this comparison, we used the Large network size for both the Inception and Residual architectures. We trained these models as described in section 2.3, with batch sizes of 128 events.

The L1 and MSE loss functions perform similarly, but the L1 loss model has a narrower error distribution across all energy ranges. Therefore, we conclude that the L1 loss function is slightly more optimal than the MSE loss function. Training with the fractional error loss produced skewed model predictions. In figure 7 we plotted the fractional reconstruction error for each evaluated model. The Frac. distribution is peaked at an overprediction of around 5%, and has a wide tail of underpredictions extending beyond -40% . The plots of the mode value and 68% quantile as a function of energy in figure 8 further show that this 5% overprediction bias and large distribution width extend throughout all true energy bins. This behavior when using the fractional loss function is not entirely understood. It may be related to the specific interplay of the network architecture, loss function, and training data.




Figure 7. Histograms of the fractional error for the InceptNet models (*left*) and the ResNet models (*right*) trained with the three different loss functions. The models trained with the L1 and MSE loss perform comparably, but the fractional (Frac.) loss models underperform.

3.3 Network input comparison

We next investigate the importance of input information concerning unresponsive wires. In order to investigate the neural networks' ability to compensate for unresponsive wires in the detector, we trained three groups of models on three different sets of information. The first group was trained on an ideal dataset with no unresponsive wires. The second group was trained on a dataset with an average of 10% unresponsive wires in each plane (URW dataset), and was given information on the locations of these wires. We marked the locations of responsive or unresponsive wires on each plane with pixel values of zero or one respectively, and recorded these on a separate set of three planes per event which were given as additional input to the neural network. The last group was trained on the same dataset as the second, but was not given any wire information.




Figure 8. The mode and absolute 68 % quantiles of the fractional reconstruction error, in each bucket of true shower energy, for the InceptNet models (*left*) and the ResNet models (*right*) trained with different loss functions.

We tried both the Inception and Residual architectures for this comparison, but used the results from sections 3.1 and 3.2 to fix the model size to Large and the loss function to L1. All models were trained on a set of 400 k images and validated on a set of 200 k, with a batch size of 128, for a total of 100 epochs, as described in section 2.3.

We then tested the performance of these three groups of models on 200 k validation events in both the ideal dataset and on the 10 % unresponsive dataset. Figure 9 the fractional error distribution on the ideal dataset for the InceptNet and ResNet models of the three groups. All the models output reconstructed energies which are within 2 % of the true values for greater than 68 % of all test events. However, the models trained on the unresponsive wire datasets slightly overpredict the energy. 49.8 % of the ideal InceptNet reconstructions are greater than the true energy, but 68.9 % (62.3 %) of the informed (uninformed) unresponsive wire InceptNet reconstructions are greater than the true energy. Similarly, 49.9 % of the ideal ResNet reconstructions are greater than the true energy, but 71.2 % (71.4 %) of the informed (uninformed) unresponsive wire ResNet reconstructions respectively are greater than the true energy.

Figure 10 shows the mode and 68% quantile of each network on the ideal dataset as a function of true energy. The models trained on the unresponsive wire dataset are trained to compensate for missing charge, so on an ideal dataset they overpredict slightly across all shower energies. We

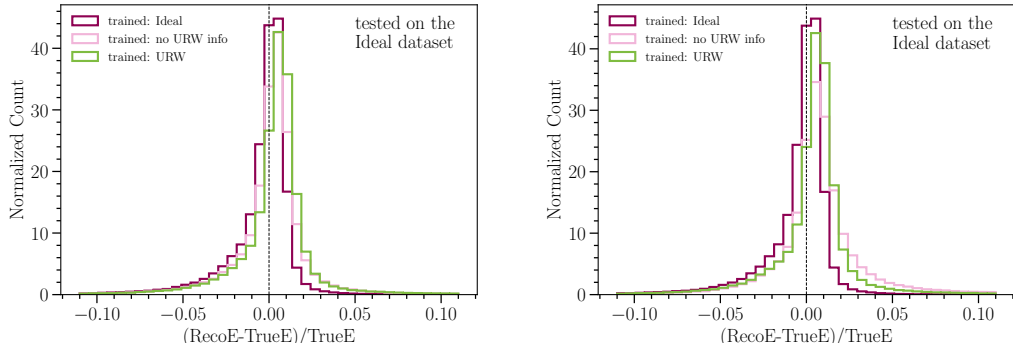


Figure 9. The fractional error of the InceptNet models (*left*) and ResNet models (*right*) trained on three different input datasets, and evaluated on the validation dataset with ideal, no unresponsive wire images.

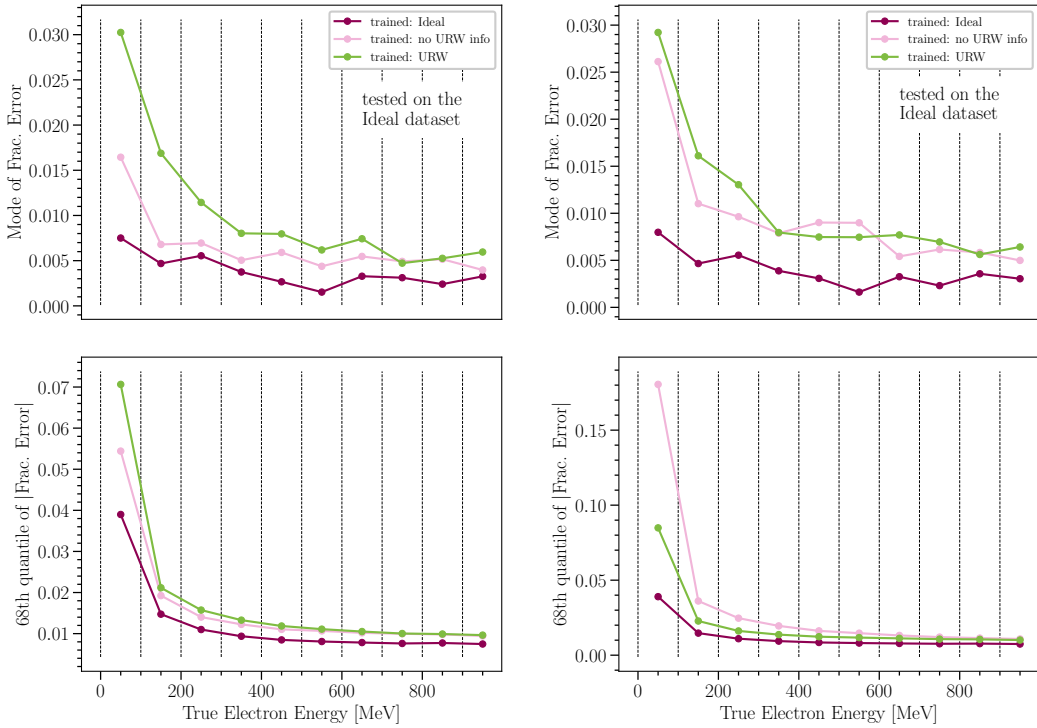


Figure 10. The mode (*above*) and absolute 68 % quantiles (*below*) of the fractional reconstruction error, in each bin of true shower energy, for the InceptNet models (*left*) and the ResNet models (*right*). These models were trained on three different input datasets, and then tested on ideal events.

expected this effect to be less significant in the model given wire location information as input, since this information could regulate the compensation for missing charge. For the ResNet models, this appears to be true: the 68% quantiles of the model trained without wire information are narrower than those of the model trained with wire information across all shower energies. This is not clear for the InceptNet models: the model trained without performs similarly, or slightly better than, the InceptNet model trained with information.

Having considered the ideal case, we now turn to the performance of these networks on shower events with unresponsive wires. Figure 11, like figure 9, shows the fractional error distribution for the InceptNet and ResNet models of the three groups described above, considering now the URW dataset instead of the ideal dataset. Additionally, we include a fourth class of models which were trained with URW information, but tested without that information (i.e., the wire status input layers are effectively null, incorrectly suggesting to the network that no dead wires exist in the image). The Ideal models do not reconstruct this dataset well; they underpredict the true energy by more than 20% in over 18% of events. The models trained with URW information but tested without also do not reconstruct well, and underpredict the true energy by more than 20% in 14.9% and 16.6% of events (for InceptNet and ResNet, correspondingly). This confirms that these models are taking advantage of the URW information given to them.

Figure 12, like figure 10, shows the mode and 68% quantile of each network as a function of true energy but again considering the URW dataset instead of the ideal dataset. Both the InceptNet and ResNet models trained and tested with URW information reconstruct shower energies with much better accuracy than the other models across all shower energies. They reconstructed 68% of the test events to within 3.1% and 3.4%, respectively, and reconstructed 95% of events to within 14.0% and 16.0% accuracy, respectively. This seems to come at the price of a slight tendency to over-estimate shower energies, especially for low energy showers, as shown by the modes in figure 12. However, the relative increase in the reconstruction accuracy compared to networks trained without URW information significantly outweighs the apparent over-prediction bias.

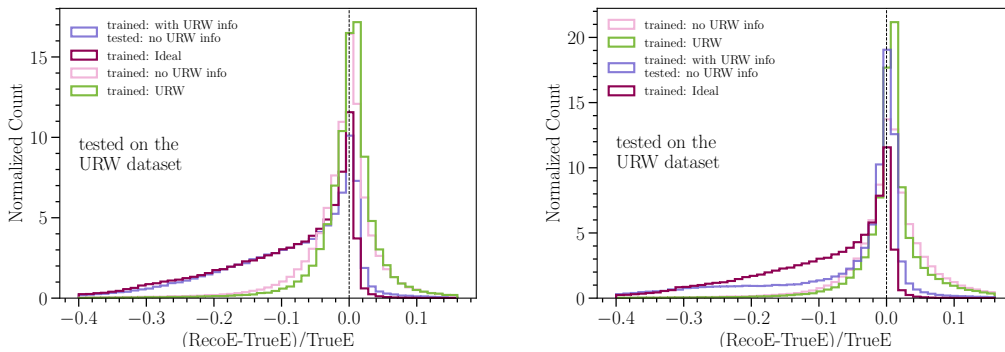


Figure 11. The fractional error of the InceptNet models (*left*) and ResNet models (*right*) trained on three different input datasets, and evaluated on the validation dataset with unresponsive wires (URW). The performance of the model trained with information on the URW locations, but tested without this information, is plotted in purple-blue.




Figure 12. The mode (*above*) and absolute 68 % quantiles (*below*) of the fractional reconstruction error, in each bucket of true shower energy, for the InceptNet models (*left*) and the ResNet models (*right*). These models were trained on three different input datasets, and then tested on URW events.

3.4 Comparison to the linear algorithm’s performance

Sections 3.1, 3.2, and 3.3 narrowed down the best-performing CNN models. In general, the “Large” networks trained under an L1 loss function and provided unresponsive wire information appear to perform optimally, especially on imperfect datasets. We now compare these candidates, implemented in both the InceptNet and ResNet architectures, directly to the linear algorithm. In addition to this comparison between the linear algorithm and neural networks in general, the tests performed in this section are further used to distinguish the performance of the InceptNet and ResNet individually.

In figure 13, we compare the performance on ideal data of the InceptNet and ResNet models trained on URW data and the linear algorithm fit to the same URW data. The linear algorithm, even with suboptimal fit parameters, still produces an excellent reconstruction: the reconstructed energy is overpredicted in 52.7 % of events, and 68 % of events have a reconstruction error within 1.99 %. The neural networks, trained on data with unresponsive wires, slightly over-estimate the energies of these ideal showers: the reconstructed energy is overpredicted by the InceptNet and ResNet models in 68.9 % and 71.2 % of events respectively. However, because the neural network fractional error distributions are narrowly peaked, 68 % of events are reconstructed to within 1.43 % and 1.52 %, respectively.

In most energy bins, the neural networks are able to match the linear algorithm's performance, producing similarly sized confidence intervals and slightly larger modes. In the lowest energy bins, the neural networks actually produce more accurate results, with a mode closer to zero and a smaller confidence interval. While the linear algorithm's calibration procedure sacrifices the fit in the lowest energy bins to secure a good fit at the rest, the neural networks' more flexible reconstruction produces a better compromise overall.




Figure 13. *Left:* the fractional reconstruction error histograms of the linear algorithm and the two CNNs, tested on the ideal validation dataset. The CNNs were Large sized and trained using the L1 loss function on data with unresponsive wires. The correction for missing charge due to unresponsive wires may be responsible for the CNN's overprediction bias on these ideal showers. *Right:* the mode (*above*) and absolute 68 % quantiles (*below*) of the fractional reconstruction error for these same two neural networks and the linear algorithm.

For the more realistic dataset with unresponsive wires, the neural networks and the linear algorithm offer different advantages. The performance of each on this dataset is compared in figure 14. The fractional error distributions for the InceptNet and ResNet models are still slightly asymmetric: 66.4 % and 56.9 % of events are overpredicted, respectively. The fractional error distribution for the linear algorithm has a long tail of underpredictions, over-estimating the energy of only 32.4 % of events (i.e. under-estimating the energy of 67.6 %). The linear algorithm can provide very accurate reconstructions for a larger fraction of events: it reconstructs 18.4 % of the total dataset to within 0.5 % error, while InceptNet and ResNet models manage 15.9 % and 15.5 % of the total dataset, respectively. By the 1 % error cutoff, the neural networks just barely overpass the linear algorithm: the linear algorithm reconstructs 29.6%, the InceptNet model 30.8%, and the ResNet model 32.2% of the total dataset to within 1 % accuracy.

When considering reconstruction on the bulk of events, however, the neural networks are the better choice. The InceptNet and ResNet models reconstructed 68 % of the test dataset to within 3.40 % and 3.15 % accuracy, respectively. The linear algorithm's interval is more than double that: 68 % of events are contained within 8.9 % fractional error. For 95 % of events, the fractional errors of the InceptNet and ResNet reconstructions are within 14.91 % and 18.49 %, respectively, while the same bound for




Figure 14. *Left:* the fractional reconstruction error histograms of the linear algorithm and the two CNNs, tested on the validation dataset with unresponsive wires. The CNNs were Large sized and trained using the L1 loss function on data with unresponsive wires. *Right:* the mode (*above*) and absolute 68 % quantiles (*below*) of the fractional reconstruction error for these same two neural networks and the linear algorithm. All three plots demonstrate the CNNs’ robustness to unresponsive wire effects.

the linear algorithm reconstruction is 53 %. When the linear algorithm fails to reconstruct an event, presumably because it passed through a dense region of unresponsive wires, it fails thoroughly.




Figure 15. The absolute 68 % quantile of the fractional error, plotted as a function of the percentage of unresponsive wires affecting the Y plane of the input images.

10 % unresponsive wires have energy resolution $> 10 \%$ in the linear algorithm, while the energy resolution of the neural network models is considerable better in this regime. One can also see that InceptNet appears to be more robust to larger URW fractions than ResNet. In figure 16 we plot the percentage of events reconstructed to within 1 % and 5 % accuracy. Interestingly, for a URW

For one-third of our test dataset, the electron showers don’t pass through any unresponsive wires in the Y-plane. On these events, the linear algorithm is extremely accurate. On the remaining two-thirds of events, however, its performance declines much more significantly than that of the neural networks. Figures 15 and 16 show this effect from two different points of view. In figure 15 we show the 68 % quantile of the absolute fractional error as a function of the percentage of unresponsive wires in an event. Events with more than 10 % unresponsive wires are very poorly reconstructed by the linear algorithm, and contribute heavily to the long tail of the fractional error histogram. From figure 15, one can see that events with at least

fraction $\lesssim 5\%$, the linear algorithm and the ResNet algorithm reconstruct events to 1% accuracy than the InceptNet algorithm. This trend reverses for events with $> 5\%$ unresponsive wires: the InceptNet algorithm reconstructs more of these events to 1% accuracy than the ResNet and linear algorithms, though the ResNet algorithm still outperforms the linear algorithm in this regime. When the error interval expands from 1% to 5%, both neural network architectures outperform the linear algorithm, even on events with fewer than 5% unresponsive wires in the Y-plane. As the URW fraction increases, the neural networks are able to reconstruct a significantly higher fraction of events to within 5% accuracy compared with the linear algorithm. Thus, one can conclude that shower energy reconstruction algorithms which use CNNs are considerably more robust to showers passing with a large fraction of unresponsive wires compared to traditional linear calibration algorithms.

Additionally, the same trend is observed in the 5% accuracy case when comparing the two network architectures: the InceptNet algorithm outperforms the ResNet algorithm for large URW fractions, while the ResNet algorithm performs better in the very small URW fraction regime ($\lesssim 5\%$ of wires). As the ResNet algorithm behaves more similarly than the InceptNet algorithm to the linear algorithm, it may be striking a balance between peak precision and bulk accuracy.

Therefore, the optimal network architecture for reconstructing shower energies in a LArTPC depends in general on the goal of the analysis and more specifically on the anticipated fraction of unresponsive wires through which the showers travel.

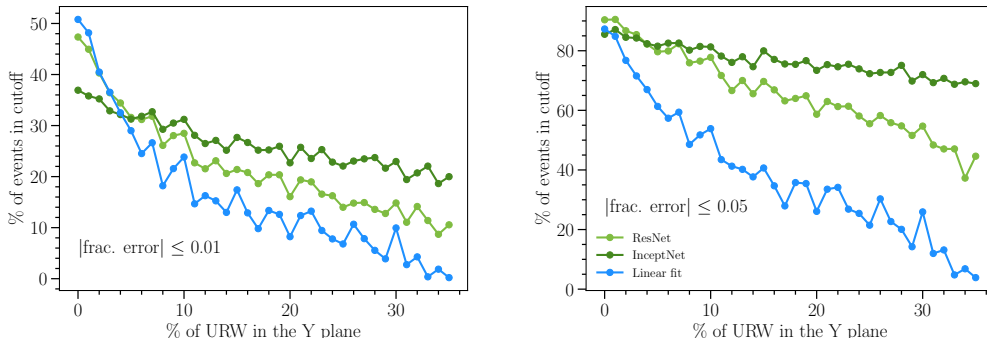


Figure 16. The percentage of events reconstructed to within 1% (*left*) or 5% (*right*) accuracy, as a function of the percentage of unresponsive wires affecting the Y plane of the input images. The linear algorithm reconstructs a larger fraction of almost-ideal events within the stricter 1% cutoff than the CNNs, but a smaller fraction of these events within the wider 5% cutoff.

4 Conclusion

In this report, we studied the ability of CNN-based algorithms to reconstruct electromagnetic shower energies in a LArTPC. Different classes of CNN algorithms were compared in their ability to reconstruct showers accurately and with minimal bias across a range of energies. These studies indicate that CNN-based shower energy reconstruction algorithms in a LArTPC are robust against showers that pass through multiple chunks of unresponsive wires in the detector. Their

competitiveness with typical clustering linear algorithms owes to the ability to reconstruct a larger fraction of events with moderate amounts of missing charge to within reasonable ($\lesssim 5\%$) accuracy. We found that CNNs could and did utilize information on the location of unresponsive wires to further improve reconstruction efficiency. Additionally, we found that the InceptNet architecture is slightly more robust than the ResNet architecture to showers that pass through larger fractions of unresponsive wires, though the latter performs better than the former in cases of few unresponsive wires.

The performance of the clustering linear algorithm in these results was in one sense optimal: since our simulated events had no background from non-shower charge depositions, there was no inefficiency due to charge being miscategorised as shower or not. We therefore expect that on more realistic datasets, the relative performance of the CNNs could further improve.

Larger CNNs, with more parameters, reconstructed the majority of events to within smaller error. In addition, the loss curves over the training epochs of all our CNNs indicate that they could still achieve slight further improvements with additional training time. Thus, the optimal performance of the CNNs is constrained by cost considerations in memory and time. Still, the main takeaway from this study is clear: CNN-based shower energy reconstruction algorithms in LArTPCs show significant improvement over traditional linear reconstruction algorithms on showers that pass through unresponsive regions of the detector.

Acknowledgments

NSF grant PHY-1801996 supported KC, JMC, AS and NWK for this work. Additionally, KC is supported by the Mariana Polonsky Slocum (1955) Memorial Fund. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302. We thank Kazuhiro Terao and John Hardin for useful discussions.

A Architecture details

Figures 17 and 18 schematically depict the convolutional layers of the Residual and Inception networks at each scale.

B Further detail on the studies

In this section, we provide further quantitative detail on the different studies described throughout this paper.

We first discuss the comparison between networks of different sizes, described in section 3.1. The fractional reconstruction error distribution for each size model in table 1 is plotted in figure 19. The widths of the distributions get narrower as the models increase in size, as expected. Table 2 reports additional quantitative information on these distributions. Specifically, we provide (1) the fraction of events with over-estimated energies (expected to be $\sim 50\%$ for an unbiased reconstruction), (2) the most probable fractional error (expected to be $\sim 0\%$ for an unbiased estimator), and (3) the 68% quantile for each fractional error distribution (i.e., the 1σ uncertainty on the energy estimation).

We next consider the loss function comparison described in section 3.2. Additional quantitative information concerning the fractional error distributions shown in figure 7 is provided in table 3.

Residual Layers, by Arch. Size:

<p>Small: Three stacked residual blocks each a single layer. Spatial dimensions are halved (downsampling) in the second and third blocks; feature size doubles in each block.</p> <p>Feature size: Down-sampling</p> <table border="1"> <tr><td>16</td><td>No</td></tr> <tr><td>32</td><td>Yes</td></tr> <tr><td>64</td><td>Yes</td></tr> </table>	16	No	32	Yes	64	Yes	<p>Medium: Four layers of two stacked blocks.</p> <table border="1"> <tr><td>16</td><td>No</td></tr> <tr><td>16</td><td>No</td></tr> <tr><td>32</td><td>Yes</td></tr> <tr><td>32</td><td>No</td></tr> <tr><td>64</td><td>Yes</td></tr> <tr><td>64</td><td>No</td></tr> <tr><td>128</td><td>Yes</td></tr> <tr><td>128</td><td>No</td></tr> </table>	16	No	16	No	32	Yes	32	No	64	Yes	64	No	128	Yes	128	No	<p>Large: Five layers of three stacked blocks.</p> <table border="1"> <tr><td>16</td><td>No</td></tr> <tr><td>16</td><td>No</td></tr> <tr><td>16</td><td>No</td></tr> <tr><td>32</td><td>Yes</td></tr> <tr><td>32</td><td>No</td></tr> <tr><td>32</td><td>No</td></tr> <tr><td>64</td><td>Yes</td></tr> <tr><td>64</td><td>No</td></tr> <tr><td>64</td><td>No</td></tr> <tr><td>128</td><td>Yes</td></tr> <tr><td>128</td><td>No</td></tr> <tr><td>128</td><td>No</td></tr> <tr><td>256</td><td>Yes</td></tr> <tr><td>256</td><td>No</td></tr> <tr><td>256</td><td>No</td></tr> </table>	16	No	16	No	16	No	32	Yes	32	No	32	No	64	Yes	64	No	64	No	128	Yes	128	No	128	No	256	Yes	256	No	256	No	<p>Huge: Six layers of four stacked blocks.</p> <table border="1"> <tr><td>16</td><td>No</td></tr> <tr><td>16</td><td>No</td></tr> <tr><td>16</td><td>No</td></tr> <tr><td>16</td><td>No</td></tr> <tr><td>32</td><td>Yes</td></tr> <tr><td>32</td><td>No</td></tr> <tr><td>32</td><td>No</td></tr> <tr><td>32</td><td>No</td></tr> <tr><td>64</td><td>Yes</td></tr> <tr><td>64</td><td>No</td></tr> <tr><td>64</td><td>No</td></tr> <tr><td>64</td><td>No</td></tr> <tr><td>128</td><td>Yes</td></tr> <tr><td>128</td><td>No</td></tr> <tr><td>128</td><td>No</td></tr> <tr><td>128</td><td>No</td></tr> <tr><td>256</td><td>Yes</td></tr> <tr><td>256</td><td>No</td></tr> <tr><td>256</td><td>No</td></tr> <tr><td>256</td><td>No</td></tr> <tr><td>512</td><td>Yes</td></tr> <tr><td>512</td><td>No</td></tr> <tr><td>512</td><td>No</td></tr> <tr><td>512</td><td>No</td></tr> </table>	16	No	16	No	16	No	16	No	32	Yes	32	No	32	No	32	No	64	Yes	64	No	64	No	64	No	128	Yes	128	No	128	No	128	No	256	Yes	256	No	256	No	256	No	512	Yes	512	No	512	No	512	No
16	No																																																																																																						
32	Yes																																																																																																						
64	Yes																																																																																																						
16	No																																																																																																						
16	No																																																																																																						
32	Yes																																																																																																						
32	No																																																																																																						
64	Yes																																																																																																						
64	No																																																																																																						
128	Yes																																																																																																						
128	No																																																																																																						
16	No																																																																																																						
16	No																																																																																																						
16	No																																																																																																						
32	Yes																																																																																																						
32	No																																																																																																						
32	No																																																																																																						
64	Yes																																																																																																						
64	No																																																																																																						
64	No																																																																																																						
128	Yes																																																																																																						
128	No																																																																																																						
128	No																																																																																																						
256	Yes																																																																																																						
256	No																																																																																																						
256	No																																																																																																						
16	No																																																																																																						
16	No																																																																																																						
16	No																																																																																																						
16	No																																																																																																						
32	Yes																																																																																																						
32	No																																																																																																						
32	No																																																																																																						
32	No																																																																																																						
64	Yes																																																																																																						
64	No																																																																																																						
64	No																																																																																																						
64	No																																																																																																						
128	Yes																																																																																																						
128	No																																																																																																						
128	No																																																																																																						
128	No																																																																																																						
256	Yes																																																																																																						
256	No																																																																																																						
256	No																																																																																																						
256	No																																																																																																						
512	Yes																																																																																																						
512	No																																																																																																						
512	No																																																																																																						
512	No																																																																																																						

Figure 17. Larger Residual Networks were designed by composing more layers out of more basic residual blocks.

Inception Layers, by Arch. Size:

<p>Small Three stacked basic inception blocks, each a single layer. Each basic block layout is the same as presented previously, although the feature dimensions at each step are adjusted to manage the overall network size. These feature dimensions are displayed below:</p> <table border="1"> <tr><td colspan="4">Max Pool (3x3, /2)</td></tr> <tr><td>16</td><td>4</td><td></td><td></td></tr> <tr><td>16</td><td>32</td><td>32</td><td>16</td></tr> <tr><td>48</td><td>52</td><td>12</td><td>16</td></tr> <tr><td>40</td><td>4</td><td></td><td></td></tr> <tr><td>64</td><td>80</td><td>16</td><td>32</td></tr> </table> <p>Downsampling occurs between each layer, using maximum pooling with a stride of 2.</p>	Max Pool (3x3, /2)				16	4			16	32	32	16	48	52	12	16	40	4			64	80	16	32	<p>Medium Three layers of two stacked inception blocks:</p> <table border="1"> <tr><td>48</td><td>8</td><td></td><td></td></tr> <tr><td>32</td><td>64</td><td>16</td><td>16</td></tr> <tr><td></td><td>64</td><td>16</td><td></td></tr> <tr><td>64</td><td>96</td><td>48</td><td>16</td></tr> <tr><td>48</td><td>8</td><td></td><td></td></tr> <tr><td>96</td><td>104</td><td>24</td><td>32</td></tr> <tr><td></td><td>64</td><td>12</td><td></td></tr> <tr><td>64</td><td>128</td><td>32</td><td>16</td></tr> <tr><td>80</td><td>16</td><td></td><td></td></tr> <tr><td>128</td><td>160</td><td>64</td><td>64</td></tr> <tr><td></td><td>96</td><td>24</td><td></td></tr> <tr><td>192</td><td>192</td><td>64</td><td>32</td></tr> </table>	48	8			32	64	16	16		64	16		64	96	48	16	48	8			96	104	24	32		64	12		64	128	32	16	80	16			128	160	64	64		96	24		192	192	64	32	<p>Large Three layers of three stacked inception blocks:</p> <table border="1"> <tr><td>96</td><td>16</td><td></td><td></td></tr> <tr><td>64</td><td>128</td><td>32</td><td>32</td></tr> <tr><td></td><td>112</td><td>24</td><td></td></tr> <tr><td>96</td><td>160</td><td>64</td><td>48</td></tr> <tr><td></td><td>128</td><td>32</td><td></td></tr> <tr><td>128</td><td>192</td><td>96</td><td>64</td></tr> <tr><td>96</td><td>16</td><td></td><td></td></tr> <tr><td>192</td><td>208</td><td>48</td><td>64</td></tr> <tr><td></td><td>112</td><td>24</td><td></td></tr> <tr><td>160</td><td>224</td><td>64</td><td>64</td></tr> <tr><td></td><td>128</td><td>24</td><td></td></tr> <tr><td>128</td><td>256</td><td>64</td><td>64</td></tr> <tr><td>160</td><td>32</td><td></td><td></td></tr> <tr><td>256</td><td>320</td><td>128</td><td>128</td></tr> <tr><td></td><td>176</td><td>40</td><td></td></tr> <tr><td>320</td><td>352</td><td>128</td><td>128</td></tr> <tr><td></td><td>192</td><td>48</td><td></td></tr> <tr><td>384</td><td>384</td><td>128</td><td>128</td></tr> </table>	96	16			64	128	32	32		112	24		96	160	64	48		128	32		128	192	96	64	96	16			192	208	48	64		112	24		160	224	64	64		128	24		128	256	64	64	160	32			256	320	128	128		176	40		320	352	128	128		192	48		384	384	128	128	<p>Huge Three layers of four stacked inception blocks:</p> <table border="1"> <tr><td>96</td><td>12</td><td></td><td></td></tr> <tr><td>152</td><td>180</td><td>52</td><td>128</td></tr> <tr><td></td><td>96</td><td>24</td><td></td></tr> <tr><td>192</td><td>224</td><td>64</td><td>160</td></tr> <tr><td></td><td>120</td><td>24</td><td></td></tr> <tr><td>216</td><td>252</td><td>72</td><td>180</td></tr> <tr><td></td><td>120</td><td>36</td><td></td></tr> <tr><td>248</td><td>292</td><td>84</td><td>208</td></tr> <tr><td>120</td><td>64</td><td></td><td></td></tr> <tr><td>128</td><td>288</td><td>128</td><td>96</td></tr> <tr><td></td><td>120</td><td>64</td><td></td></tr> <tr><td>152</td><td>348</td><td>152</td><td>116</td></tr> <tr><td></td><td>180</td><td>96</td><td></td></tr> <tr><td>180</td><td>406</td><td>180</td><td>130</td></tr> <tr><td></td><td>240</td><td>96</td><td></td></tr> <tr><td>206</td><td>460</td><td>206</td><td>152</td></tr> <tr><td>240</td><td>80</td><td></td><td></td></tr> <tr><td>116</td><td>576</td><td>288</td><td>172</td></tr> <tr><td></td><td>240</td><td>120</td><td></td></tr> <tr><td>128</td><td>640</td><td>320</td><td>192</td></tr> <tr><td></td><td>240</td><td>160</td><td></td></tr> <tr><td>140</td><td>704</td><td>352</td><td>212</td></tr> <tr><td></td><td>240</td><td>160</td><td></td></tr> <tr><td>154</td><td>768</td><td>384</td><td>230</td></tr> </table>	96	12			152	180	52	128		96	24		192	224	64	160		120	24		216	252	72	180		120	36		248	292	84	208	120	64			128	288	128	96		120	64		152	348	152	116		180	96		180	406	180	130		240	96		206	460	206	152	240	80			116	576	288	172		240	120		128	640	320	192		240	160		140	704	352	212		240	160		154	768	384	230
Max Pool (3x3, /2)																																																																																																																																																																																																																																																			
16	4																																																																																																																																																																																																																																																		
16	32	32	16																																																																																																																																																																																																																																																
48	52	12	16																																																																																																																																																																																																																																																
40	4																																																																																																																																																																																																																																																		
64	80	16	32																																																																																																																																																																																																																																																
48	8																																																																																																																																																																																																																																																		
32	64	16	16																																																																																																																																																																																																																																																
	64	16																																																																																																																																																																																																																																																	
64	96	48	16																																																																																																																																																																																																																																																
48	8																																																																																																																																																																																																																																																		
96	104	24	32																																																																																																																																																																																																																																																
	64	12																																																																																																																																																																																																																																																	
64	128	32	16																																																																																																																																																																																																																																																
80	16																																																																																																																																																																																																																																																		
128	160	64	64																																																																																																																																																																																																																																																
	96	24																																																																																																																																																																																																																																																	
192	192	64	32																																																																																																																																																																																																																																																
96	16																																																																																																																																																																																																																																																		
64	128	32	32																																																																																																																																																																																																																																																
	112	24																																																																																																																																																																																																																																																	
96	160	64	48																																																																																																																																																																																																																																																
	128	32																																																																																																																																																																																																																																																	
128	192	96	64																																																																																																																																																																																																																																																
96	16																																																																																																																																																																																																																																																		
192	208	48	64																																																																																																																																																																																																																																																
	112	24																																																																																																																																																																																																																																																	
160	224	64	64																																																																																																																																																																																																																																																
	128	24																																																																																																																																																																																																																																																	
128	256	64	64																																																																																																																																																																																																																																																
160	32																																																																																																																																																																																																																																																		
256	320	128	128																																																																																																																																																																																																																																																
	176	40																																																																																																																																																																																																																																																	
320	352	128	128																																																																																																																																																																																																																																																
	192	48																																																																																																																																																																																																																																																	
384	384	128	128																																																																																																																																																																																																																																																
96	12																																																																																																																																																																																																																																																		
152	180	52	128																																																																																																																																																																																																																																																
	96	24																																																																																																																																																																																																																																																	
192	224	64	160																																																																																																																																																																																																																																																
	120	24																																																																																																																																																																																																																																																	
216	252	72	180																																																																																																																																																																																																																																																
	120	36																																																																																																																																																																																																																																																	
248	292	84	208																																																																																																																																																																																																																																																
120	64																																																																																																																																																																																																																																																		
128	288	128	96																																																																																																																																																																																																																																																
	120	64																																																																																																																																																																																																																																																	
152	348	152	116																																																																																																																																																																																																																																																
	180	96																																																																																																																																																																																																																																																	
180	406	180	130																																																																																																																																																																																																																																																
	240	96																																																																																																																																																																																																																																																	
206	460	206	152																																																																																																																																																																																																																																																
240	80																																																																																																																																																																																																																																																		
116	576	288	172																																																																																																																																																																																																																																																
	240	120																																																																																																																																																																																																																																																	
128	640	320	192																																																																																																																																																																																																																																																
	240	160																																																																																																																																																																																																																																																	
140	704	352	212																																																																																																																																																																																																																																																
	240	160																																																																																																																																																																																																																																																	
154	768	384	230																																																																																																																																																																																																																																																

Figure 18. Larger Inception Networks were designed by adding more basic inception blocks to each layer, and by increasing the output feature sizes of all convolutional filters throughout.

This table confirms the conclusion in the main text; the model trained using an L1 loss function performs slightly better than the model trained under an MSE loss function, while the model trained with a fractional error loss function seems to significantly over-predict the energy of most showers.

In table 4 we provide additional quantitative information on the fractional error distributions in figures 9 and 11. This table again confirms the conclusions in the main text; when estimating energies of showers without URWs, the networks trained on events with URWs tend to over-predict compared to the networks trained on ideal events. However, the networks trained on images with URWs perform much better on events which also include URWs. This is especially true when information on the URWs are provided as input to the networks. It is also clear that the networks trained with input information on URW are making use of this information, as the prediction uncertainty (given by 68% quantiles in table 4) increases when this input is removed.

Finally, we report in table 5 additional quantitative information of the fractional error distributions shown in figure 13 and 14. This table considers the performance of the linear algorithm and the neural networks on evaluation datasets simulating both ideal and imperfect detectors. As discussed in section 3.4, one can see from the table that the neural network algorithms are in general more robust to unresponsive wires introduced in the imperfect dataset.




Figure 19. The normalized distributions of the fractional reconstruction error for the different sized InceptNet (*left*) and ResNet (*right*) models.

C Systematic variation study

We also tested the robustness of the CNN models and the linear algorithm to noise in the data. We applied noisy transformations to the input data, and tested the performance of the pre-trained models. We found that all models had predictable performance declines.

C.1 Flat noise

The first test we tried was applying a flat 90% or 110% multiplier to all the charge in an image. This will necessarily shift the pre-trained linear algorithm’s fractional error histogram by nearly the same amount. In theory, a CNN could adjust for this flat noise by considering the length of a shower in relation to its total charge. Additionally, in a dataset with background signals, the change in the

Table 2. Quantitative properties of the fractional error histograms in figure 19. See text for a description of the variables shown in each column.

Architecture	Size	% > 0	Mode (% error)	68 % quantile (% error)
InceptNet	Small	50.2	0.52	3.96
InceptNet	Medium	45.5	0.14	3.86
InceptNet	Large	56.9	0.71	3.15
InceptNet	Huge	47.6	0.25	2.79
ResNet	Small	50.7	0.46	3.89
ResNet	Medium	53.2	0.56	3.48
ResNet	Large	66.4	0.78	3.40
ResNet	Huge	47.1	0.32	2.56

Table 3. Quantitative properties of the fractional error histograms in figure 7.

Architecture	Loss	% > 0	Mode (% error)	68 % quantile (% error)
InceptNet	L1	56.9	0.71	3.15
InceptNet	Frac.	45.6	5.04	8.56
InceptNet	MSE	59.8	0.95	3.25
ResNet	L1	66.4	0.78	3.40
ResNet	Frac.	45.5	6.07	8.46
ResNet	MSE	61.4	0.89	4.23




Figure 20. Histograms of the fractional error on input events with a 90%, 100%, and 110% multiplier applied to the pixel charge Q , for the InceptNet (*left*) and ResNet (*right*) models, compared to the linear fit.

background charge could serve as a reference point. However, our dataset had no background, and our CNNs' performance, plotted in figure 20 suffered the same bias shifts as the linear algorithm.

Table 4. Quantitative properties of the fractional error histograms in figure 9 and figure 11

Architecture	Training Data	Testing Data	% > 0	Mode (% error)	68 % quantile (% error)
InceptNet	Ideal	Ideal	49.8	0.36	1.05
InceptNet	URW, no info	Ideal	62.3	0.55	1.36
InceptNet	URW	Ideal	68.9	0.81	1.43
ResNet	Ideal	Ideal	49.9	0.37	1.05
ResNet	URW, no info	Ideal	71.4	0.57	1.97
ResNet	URW	Ideal	71.2	0.70	1.15
InceptNet	Ideal	URW	12.7	0.11	13.33
InceptNet	URW, no info	URW	46.3	0.33	3.90
InceptNet	URW	URW	56.9	0.71	3.14
InceptNet	URW	URW, no info	20.3	0.41	12.33
ResNet	Ideal	URW	12.7	0.11	13.33
ResNet	URW, no info	URW	60.5	0.35	4.87
ResNet	URW	URW	66.4	0.78	3.40
ResNet	URW	URW, no info	33.6	0.34	7.86

Table 5. Quantitative properties of the fractional error histograms in figure 13 and 14

Architecture	Training Data	Testing Data	% > 0	Mode (% error)	68 % quantile (% error)
Linear	URW	Ideal	52.8	0.07	1.99
InceptNet	URW	Ideal	68.9	0.81	1.43
ResNet	URW	Ideal	71.2	0.70	1.52
Linear	URW	URW	32.4	0.89	8.92
InceptNet	URW	URW	56.9	0.71	3.15
ResNet	URW	URW	66.4	0.78	3.40

C.2 Gaussian per-wire noise

We also checked the models' robustness under Gaussian noise on the detector wires. For this test, we multiplied each wire (column) in the input images by a random value pulled from a Gaussian distribution centered on 1 with width 0.1, and then evaluated the trained model's performance on this noisy dataset. We expected this kind of noise to increase the performance variance. This effect is what we see in figure 21, where we plotted the resulting fractional error distribution.

On the dataset without noise, the linear algorithm predicts the top 68 % of shower energies to within 10.1 %, and on that with noise it predicts them to within 16.5 %. The neural networks perform a little better: on the dataset without and with respectively, the InceptNet model predicts 68 % of events to within 3.2 % and 11.5 %, while the ResNet model predicts them to within 3.4 % and 11.6 %.




Figure 21. Histograms of the fractional error on input events with Gaussian distributed per-wire noise, for the InceptNet (*left*) and ResNet (*right*) models, compared to the linear fit.

References

- [1] C. Rubbia, *The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors*, Tech. Rep., [CERN-EP-INT-77-08](#), [CERN-EP-77-08](#), CERN (1977).
- [2] DUNE collaboration, *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics*, [arXiv:2002.03005](#)
- [3] MicroBooNE collaboration, *Design and Construction of the MicroBooNE Detector*, [2017 JINST 12 P02017](#) [[arXiv:1612.05824](#)].
- [4] P.A. Machado, O. Palamara and D.W. Schmitz, *The Short-Baseline Neutrino Program at Fermilab*, [Ann. Rev. Nucl. Part. Sci.](#) **69** (2019) 363 [[arXiv:1903.04608](#)].
- [5] ARGONeUT collaboration, *First Measurements of Inclusive Muon Neutrino Charged Current Differential Cross Sections on Argon*, [Phys. Rev. Lett.](#) **108** (2012) 161802 [[arXiv:1111.0103](#)].
- [6] MicroBooNE collaboration, *Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber*, [2017 JINST 12 P03011](#) [[arXiv:1611.05531](#)].
- [7] MicroBooNE collaboration, *The Pandora multi-algorithm approach to automated pattern recognition of cosmic-ray muon and neutrino events in the MicroBooNE detector*, [Eur. Phys. J. C](#) **78** (2018) 82 [[arXiv:1708.03135](#)].
- [8] MicroBooNE collaboration, *Deep neural network for pixel-level electromagnetic particle identification in the MicroBooNE liquid argon time projection chamber*, [Phys. Rev. D](#) **99** (2019) 092001 [[arXiv:1808.07269](#)].
- [9] X. Qian, C. Zhang, B. Viren and M. Diwan, *Three-dimensional Imaging for Large LArTPCs*, [2018 JINST 13 P05032](#) [[arXiv:1803.04850](#)].
- [10] MicroBooNE collaboration, *Electromagnetic shower reconstruction and energy validation with Michel electrons and π^0 samples for the deep-learning-based analyses in MicroBooNE*, [2021 JINST 16 T12017](#) [[arXiv:2110.11874](#)].
- [11] DeepLearnPhysics, *Particle imaging in liquid argon (pilarnet)*, <https://osf.io/bu4fp/>, 2020.
- [12] GEANT4 collaboration, *GEANT4—a simulation toolkit*, [Nucl. Instrum. Meth. A](#) **506** (2003) 250.

- [13] C. Adams, K. Terao and T. Wongjirad, *PILArNet: Public Dataset for Particle Imaging Liquid Argon Detectors in High Energy Physics*, [arXiv:2006.01993](#).
- [14] DUNE collaboration, *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume IV: Far Detector Single-phase Technology*, 2020 *JINST* **15** T08010 [[arXiv:2002.03010](#)].
- [15] MicroBooNE collaboration, *Noise Characterization and Filtering in the MicroBooNE Liquid Argon TPC*, 2017 *JINST* **12** P08003 [[arXiv:1705.07341](#)].
- [16] K. He, X. Zhang, S. Ren and J. Sun, *Deep residual learning for image recognition*, in *IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, U.S.A., 27–30 June 2016, pp. 770–778.
- [17] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov et al., *Going Deeper with Convolutions*, [arXiv:1409.4842](#).
- [18] P. Baldi, J. Bian, L. Hertel and L. Li, *Improved Energy Reconstruction in NOvA with Regression Convolutional Neural Networks*, *Phys. Rev. D* **99** (2019) 012011 [[arXiv:1811.04557](#)].
- [19] MicroBooNE collaboration, *Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE*, *Phys. Rev. D* **103** (2021) 052012 [[arXiv:2012.08513](#)].
- [20] L. Bottou, *Online algorithms and stochastic approximations*, in *Online Learning and Neural Networks*, D. Saad, ed., Cambridge University Press, Cambridge, U.K. (1998), <http://leon.bottou.org/papers/bottou-98x>.
- [21] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan et al., *Pytorch: An imperative style, high-performance deep learning library*, in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox and R. Garnett, eds., pp. 8024–8035, Curran Associates, Inc. (2019), <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.

Appendix B

Signal Event Displays from the Two-Body CCQE Analysis

Figure 5-1 in chapter 5 shows the event display for one of the 25 ν_e -candidate selected events in the $1e1p$ dataset. We provide the remaining 24 event displays in figure B-1 to figure B-24, in ascending order in energy. The top row of each display shows the pixel intensity image, while the bottom row shows the SparseSSNet pixel labels. From left to right, the columns correspond to the U,V, and Y plane images.

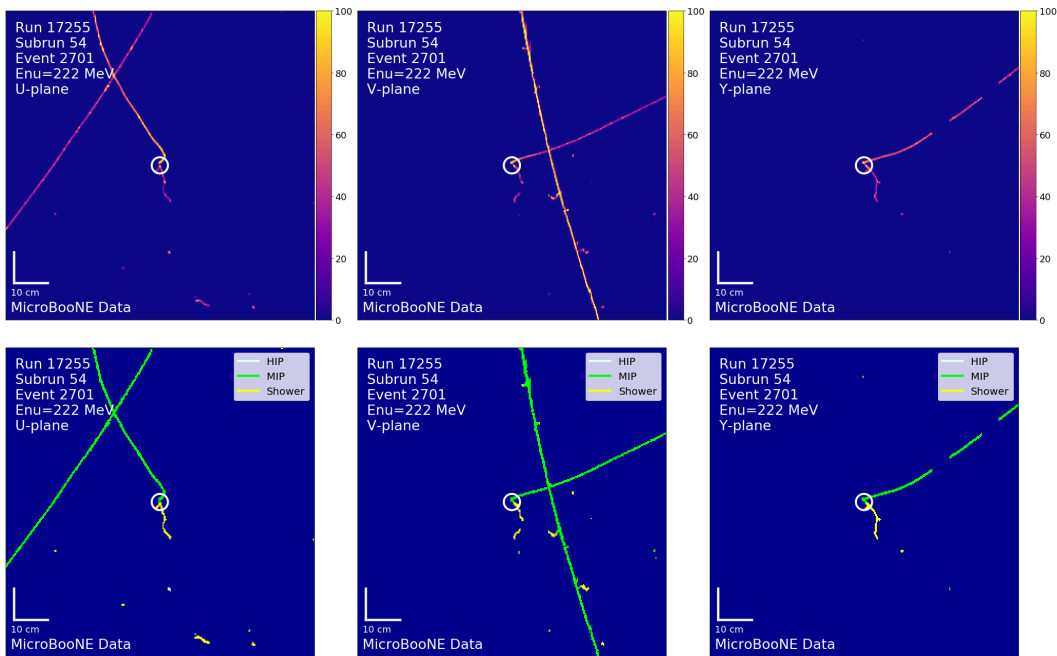


Figure B-1: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

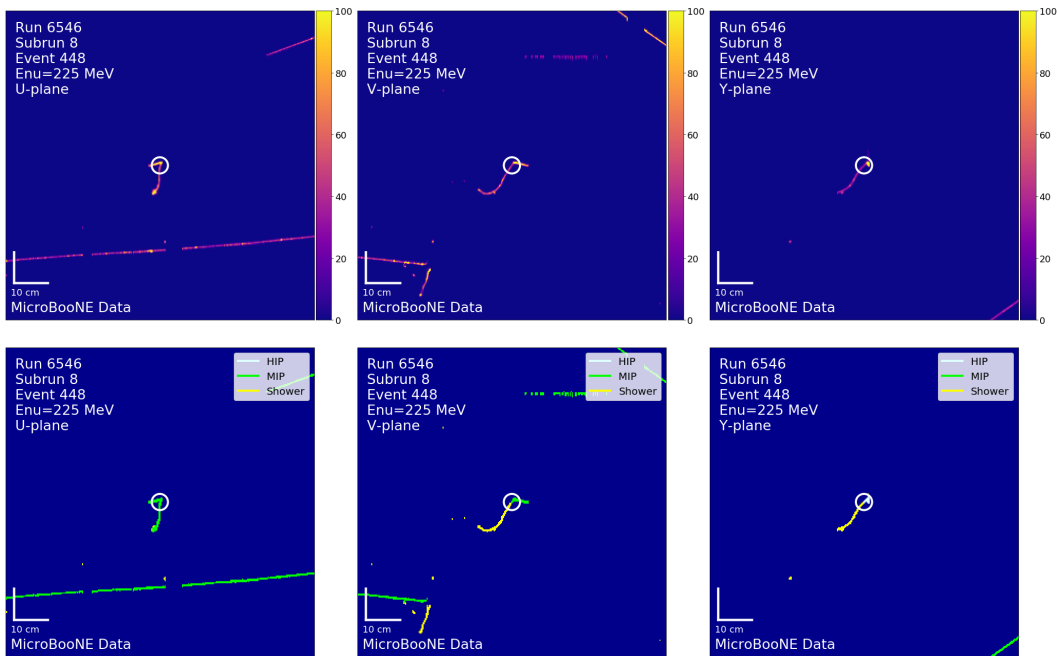


Figure B-2: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

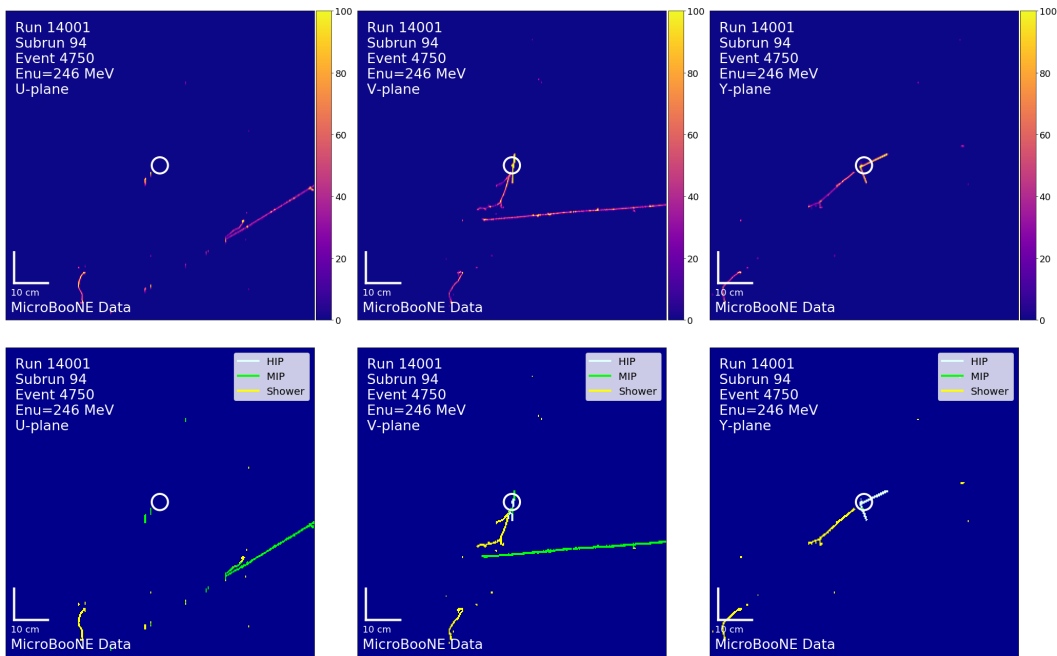


Figure B-3: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

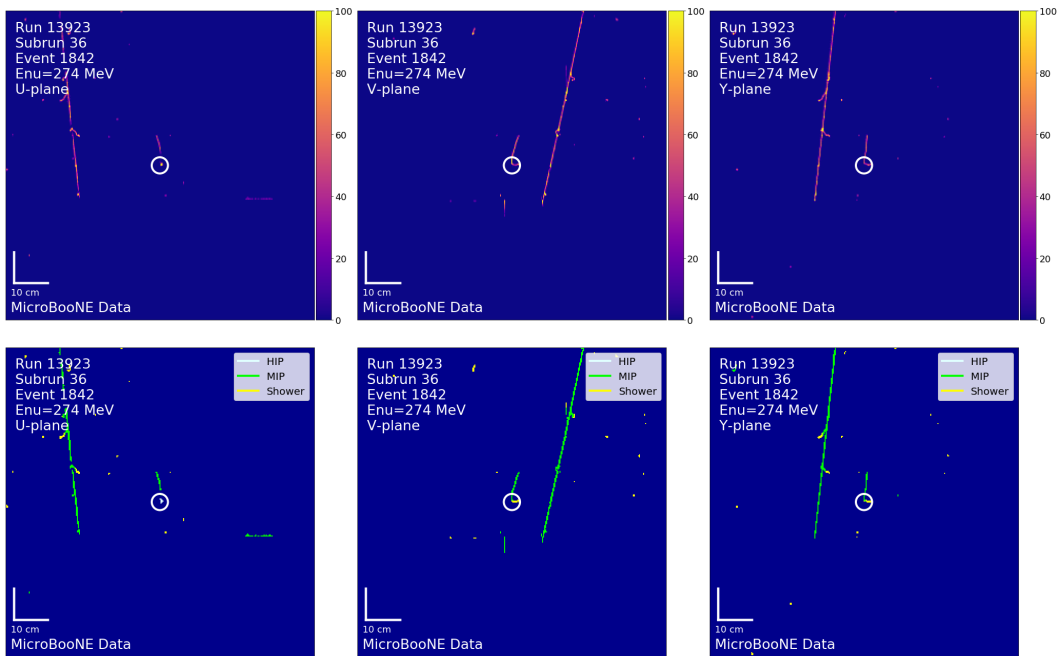


Figure B-4: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

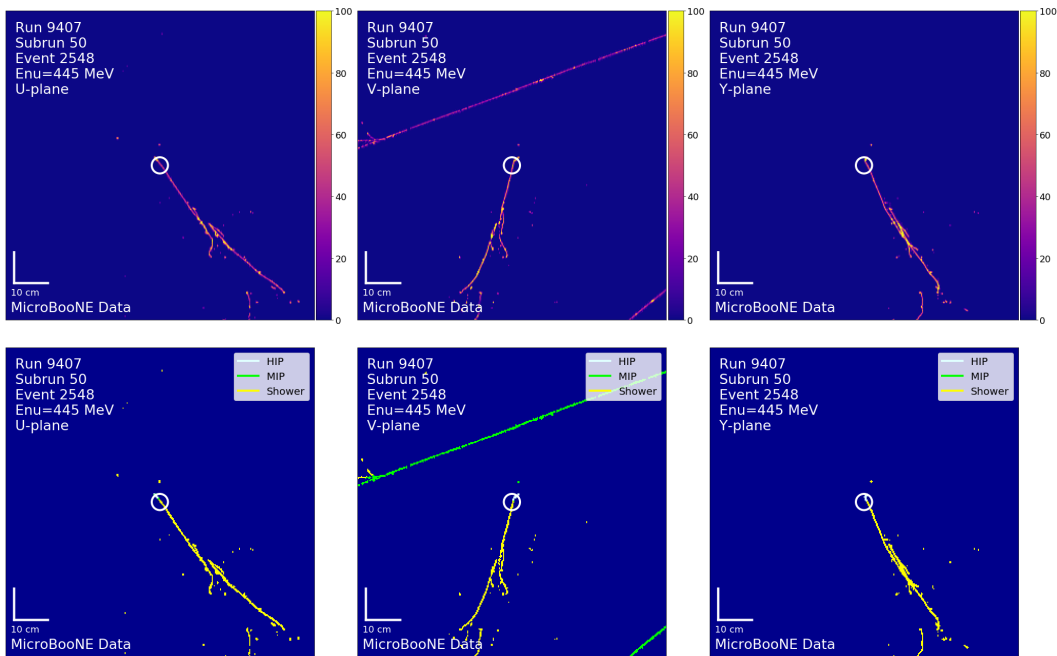


Figure B-5: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

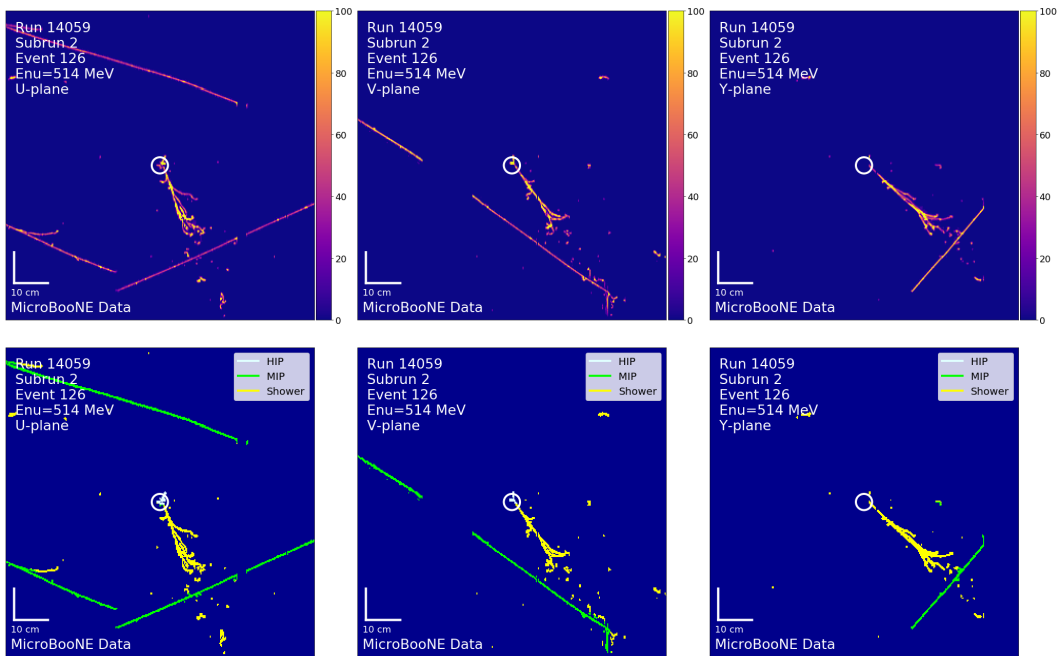


Figure B-6: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

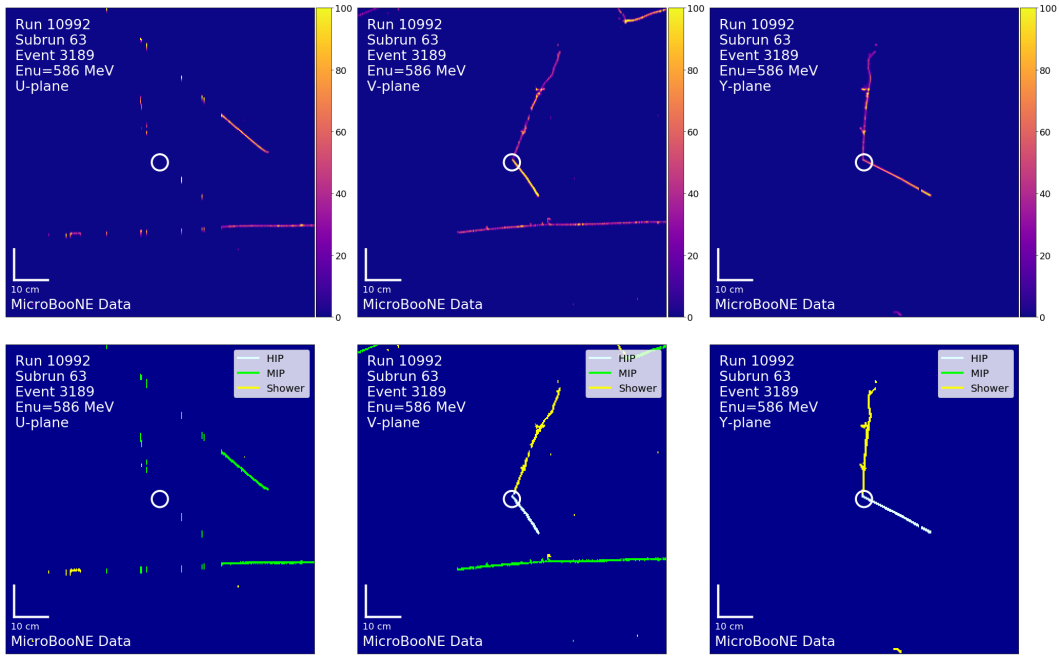


Figure B-7: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

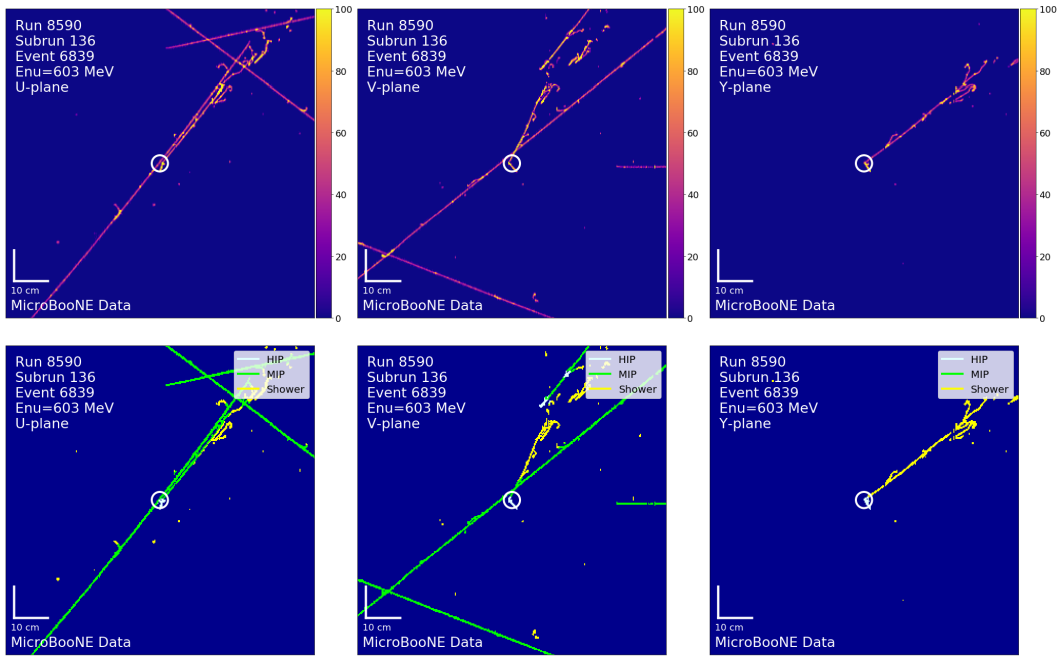


Figure B-8: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

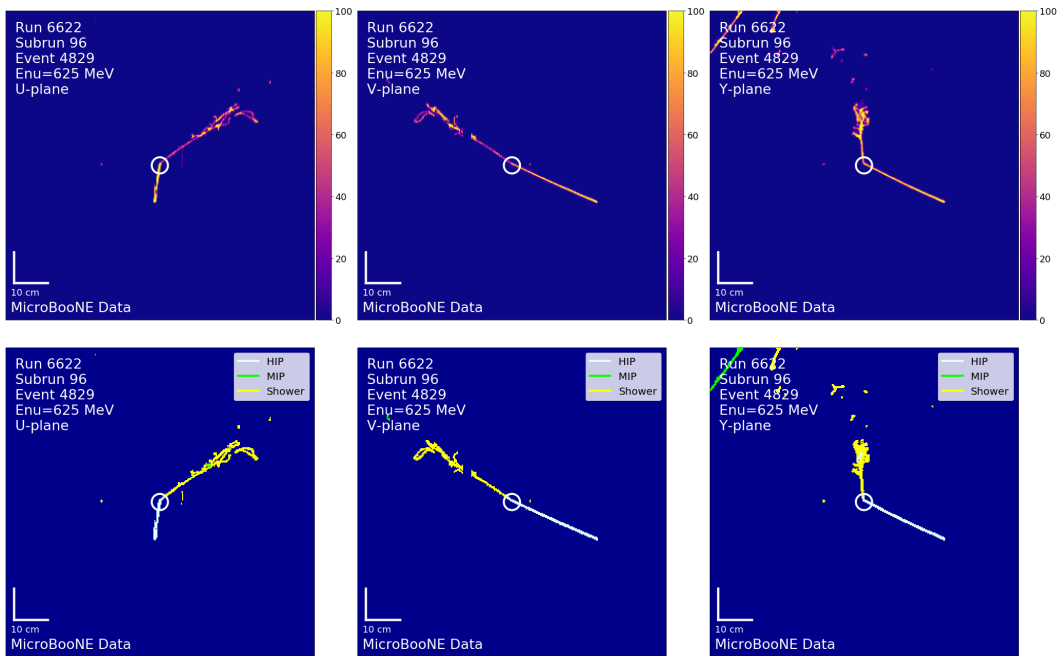


Figure B-9: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

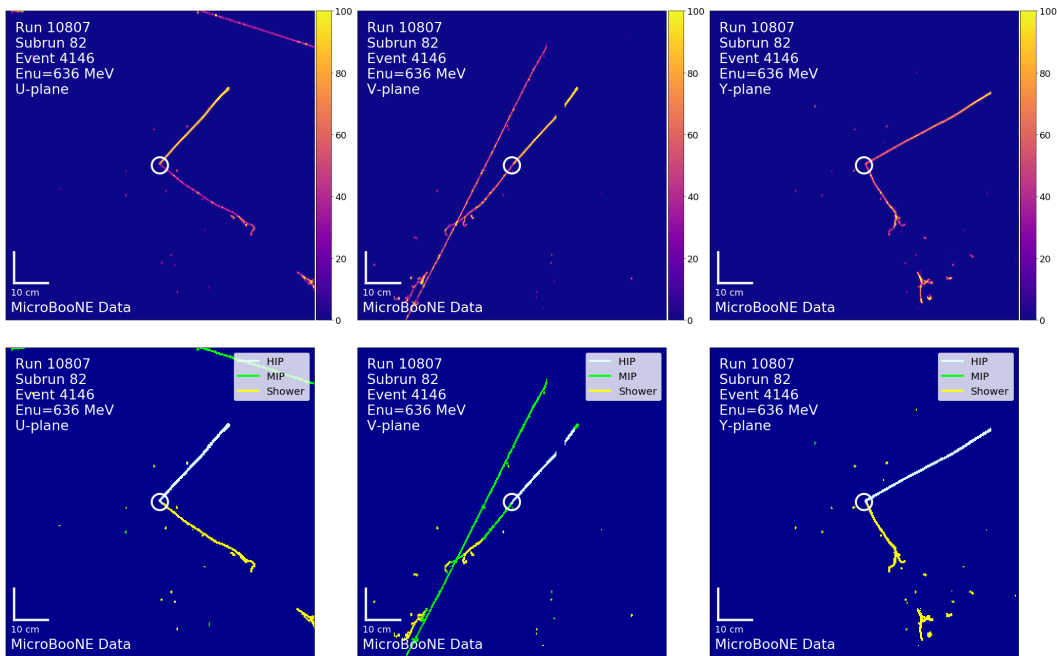


Figure B-10: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

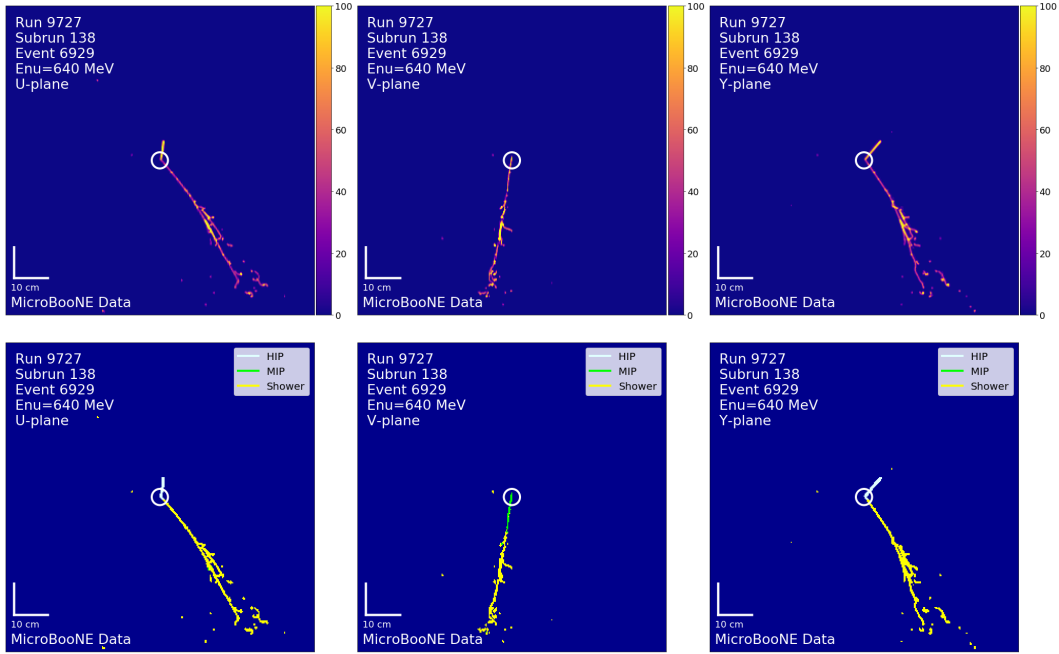


Figure B-11: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

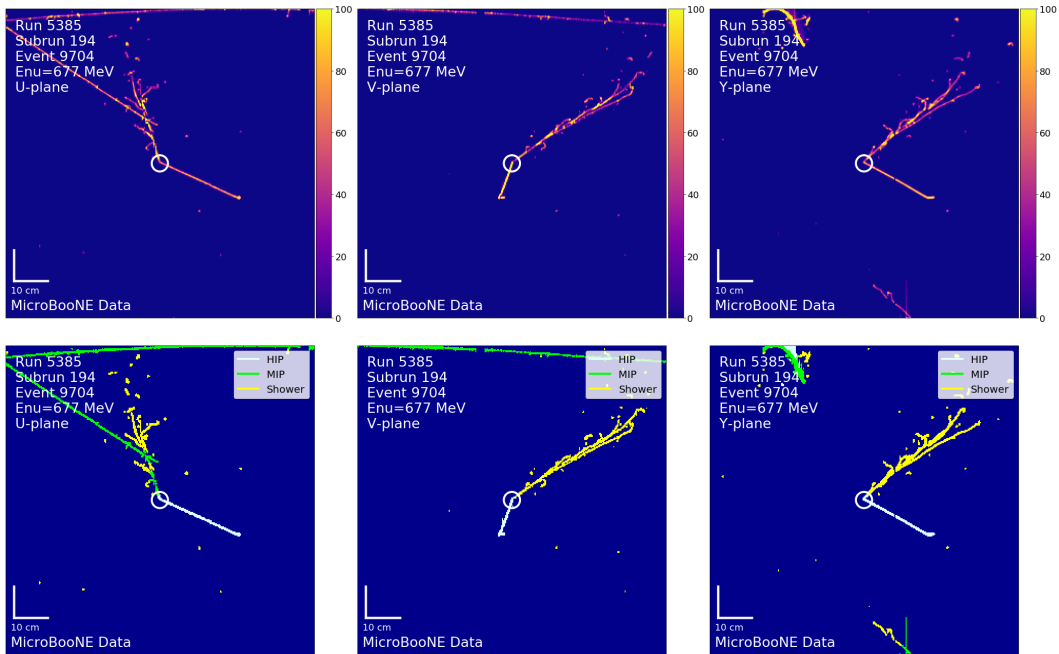


Figure B-12: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

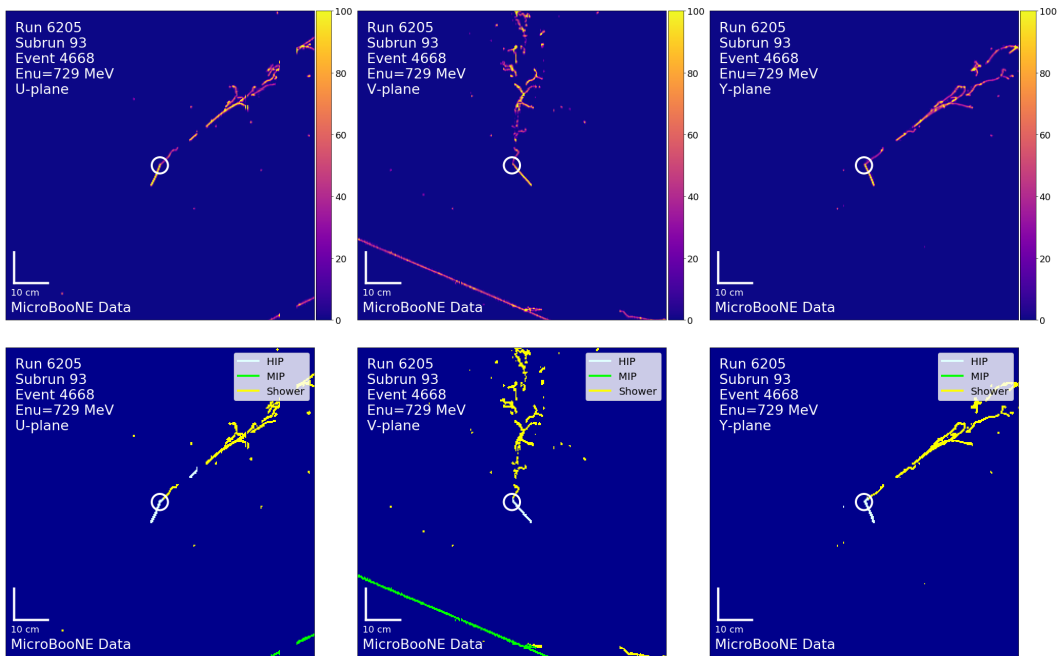


Figure B-13: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

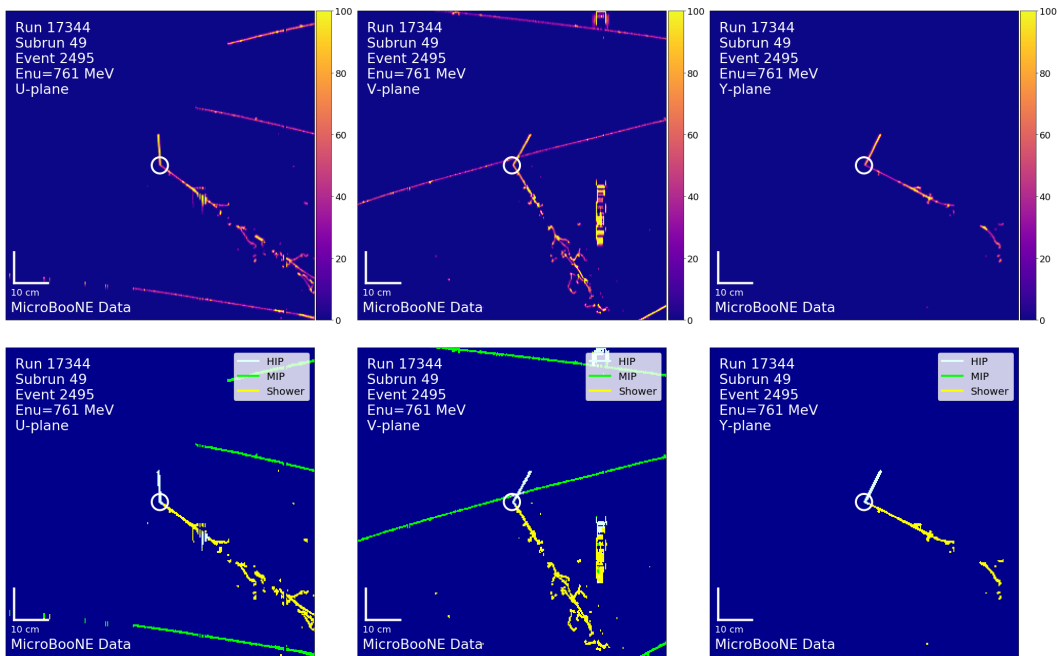


Figure B-14: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

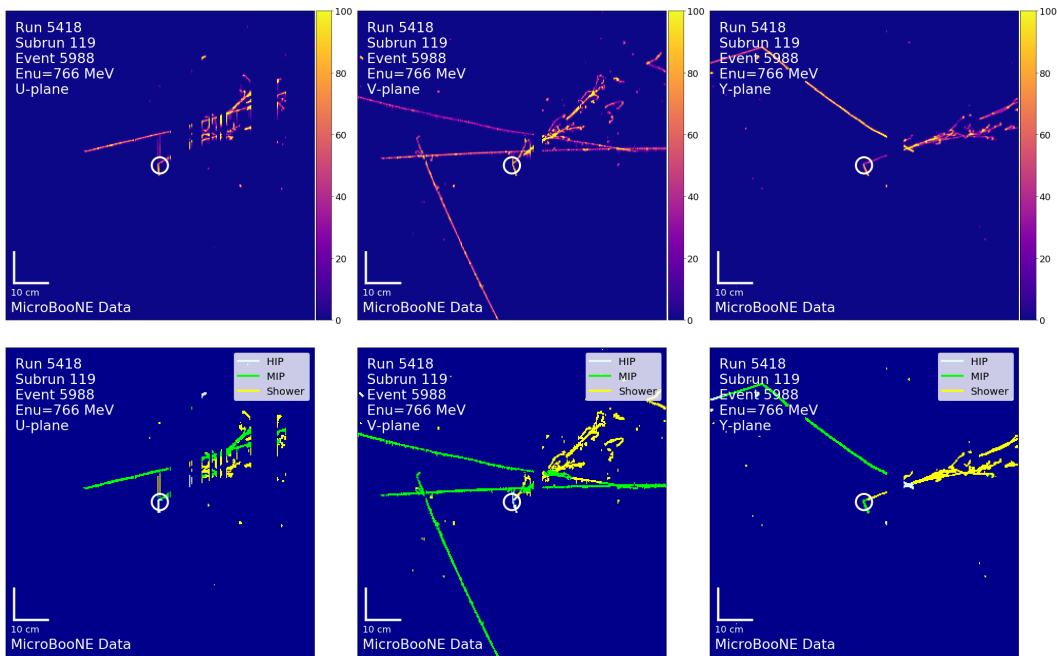


Figure B-15: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

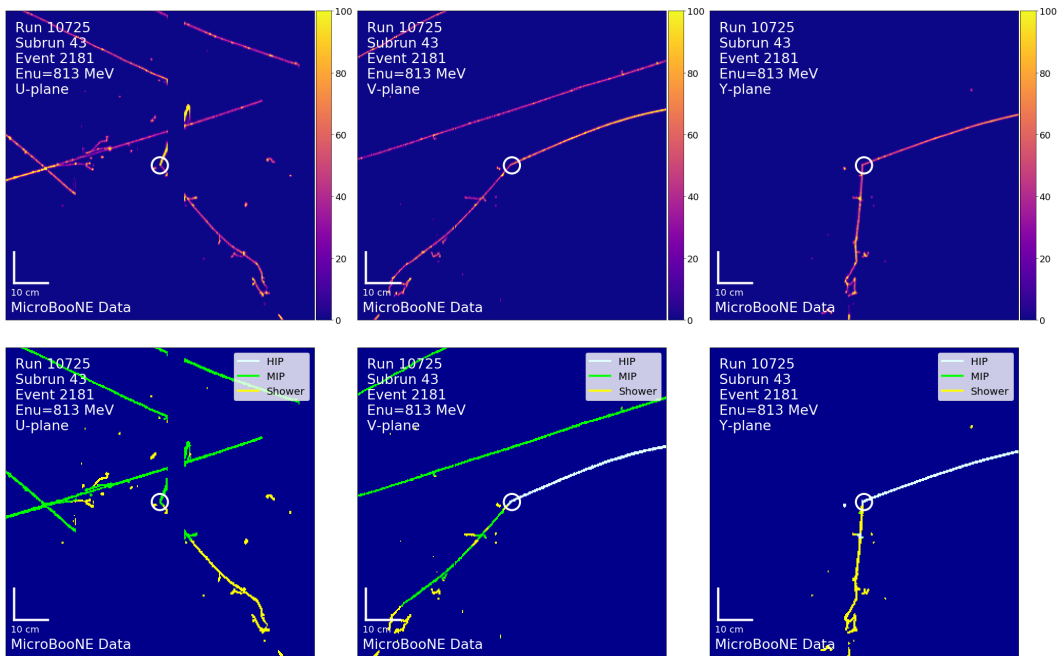


Figure B-16: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.




Figure B-17: Top: pixel intensity; Bottom: `SparseSSNet` labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

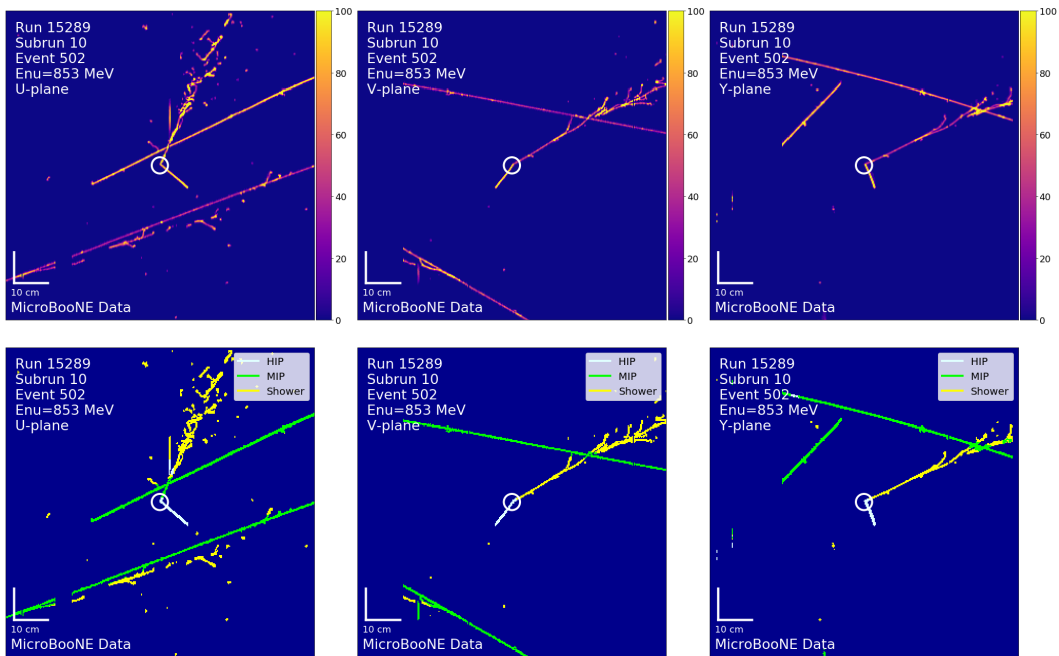


Figure B-18: Top: pixel intensity; Bottom: `SparseSSNet` labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

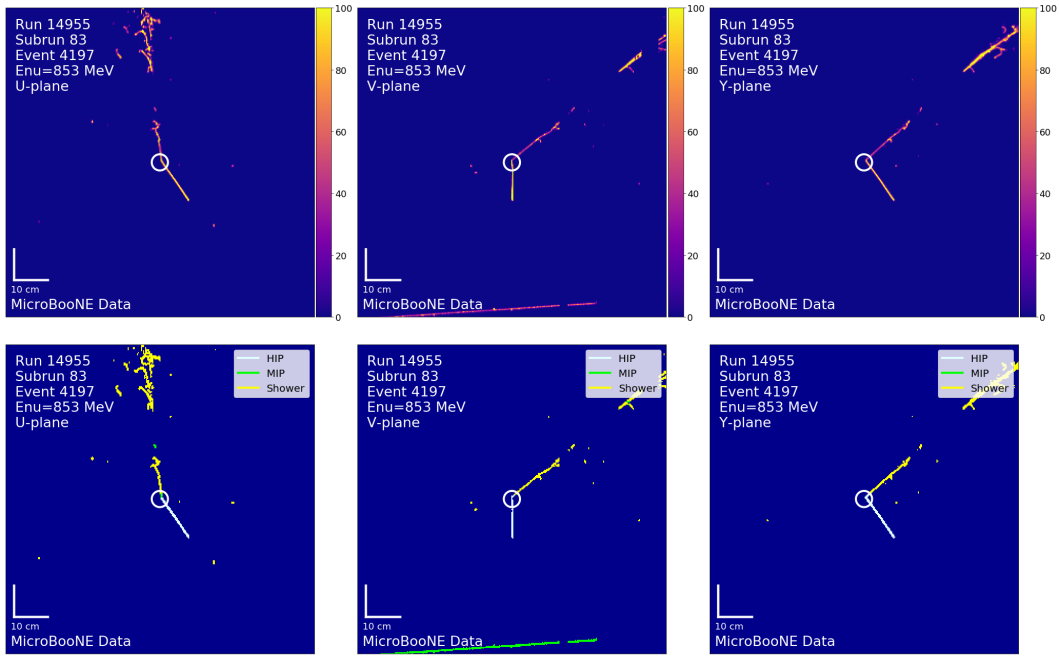


Figure B-19: Top: pixel intensity; Bottom: **SparseSSNet** labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

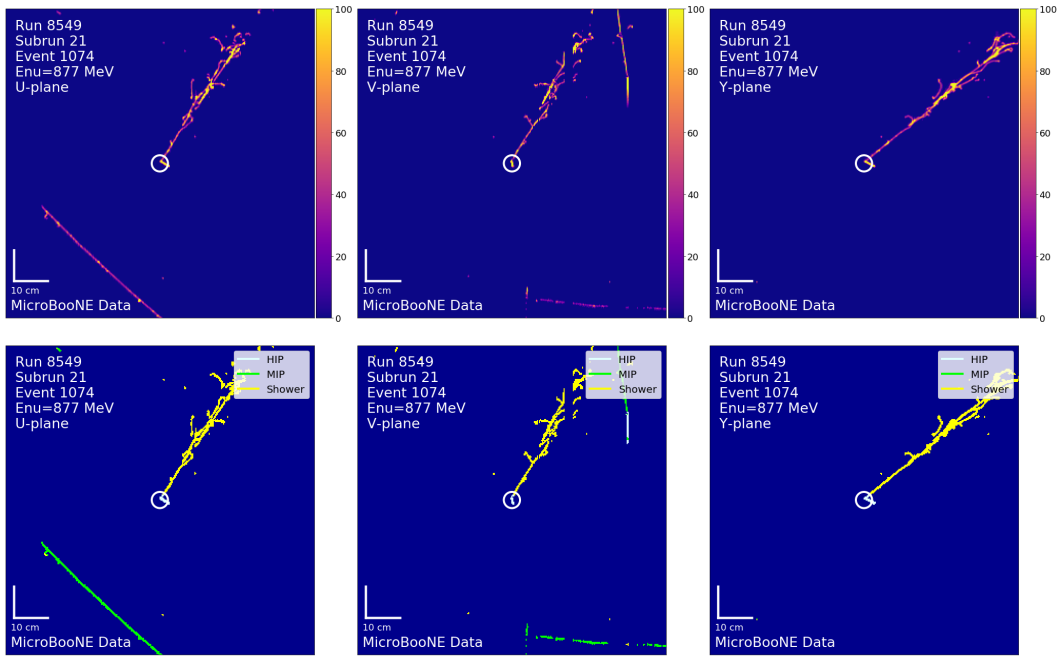


Figure B-20: Top: pixel intensity; Bottom: **SparseSSNet** labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

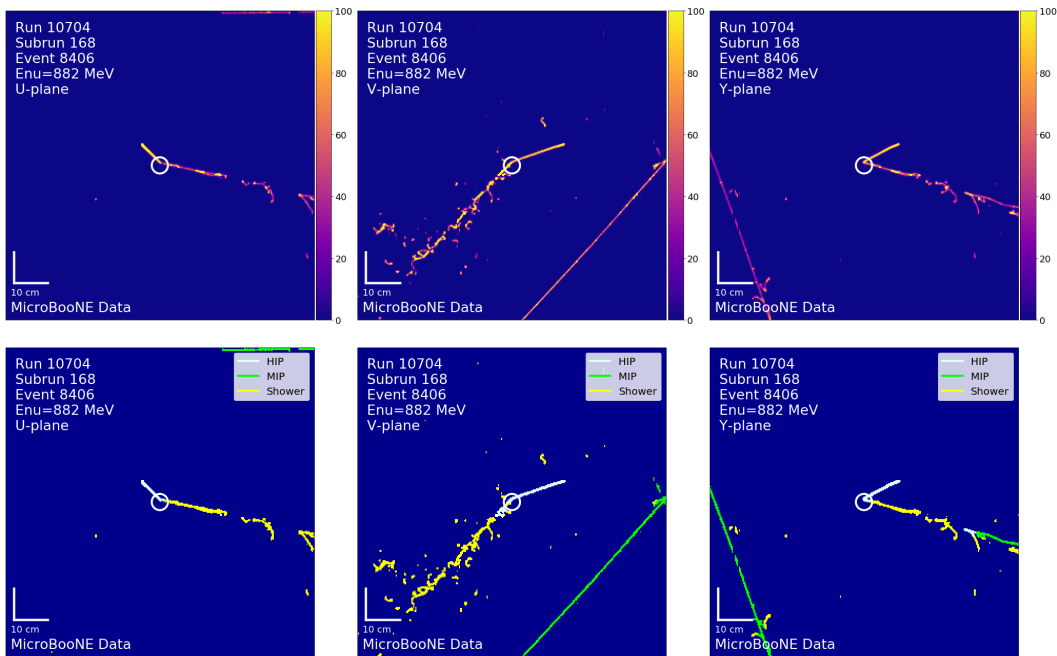


Figure B-21: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

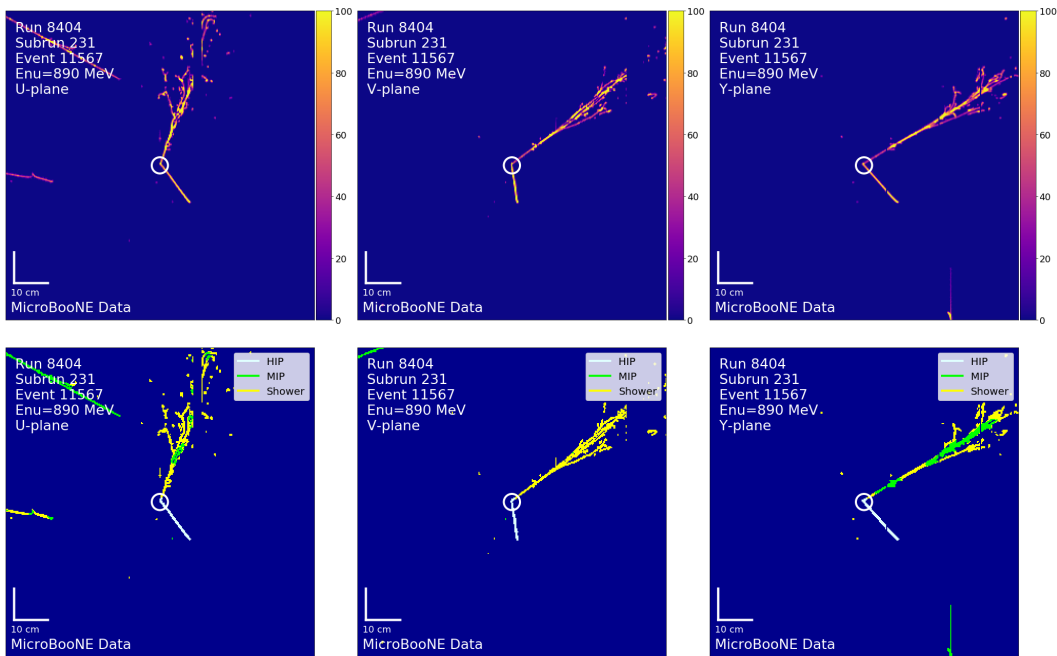


Figure B-22: Top: pixel intensity; Bottom: SparseSSNet labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

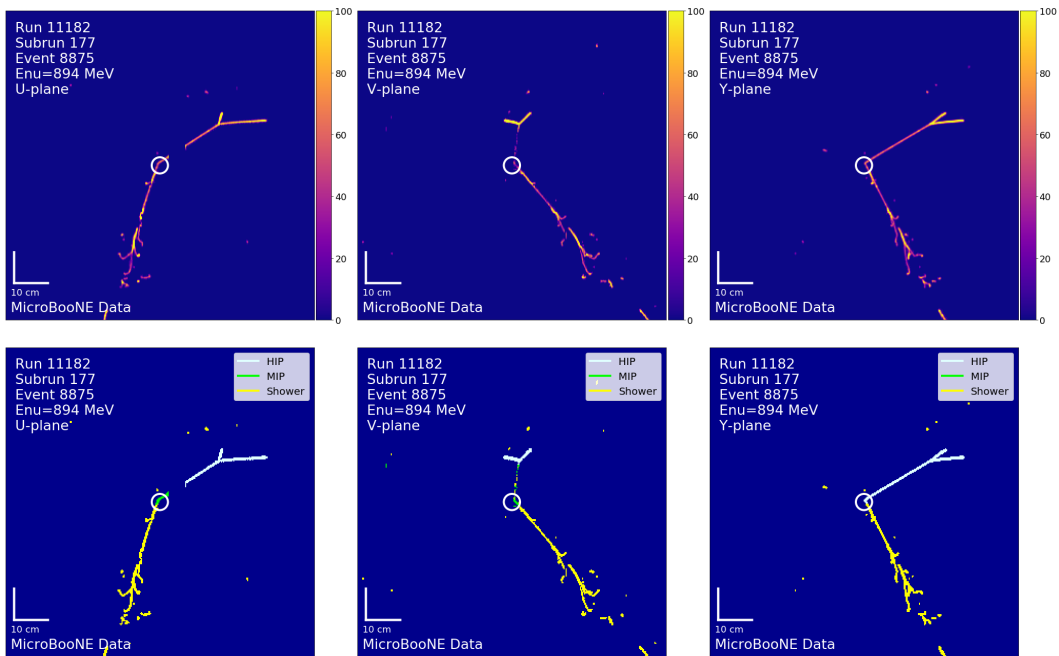


Figure B-23: Top: pixel intensity; Bottom: `SparseSSNet` labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

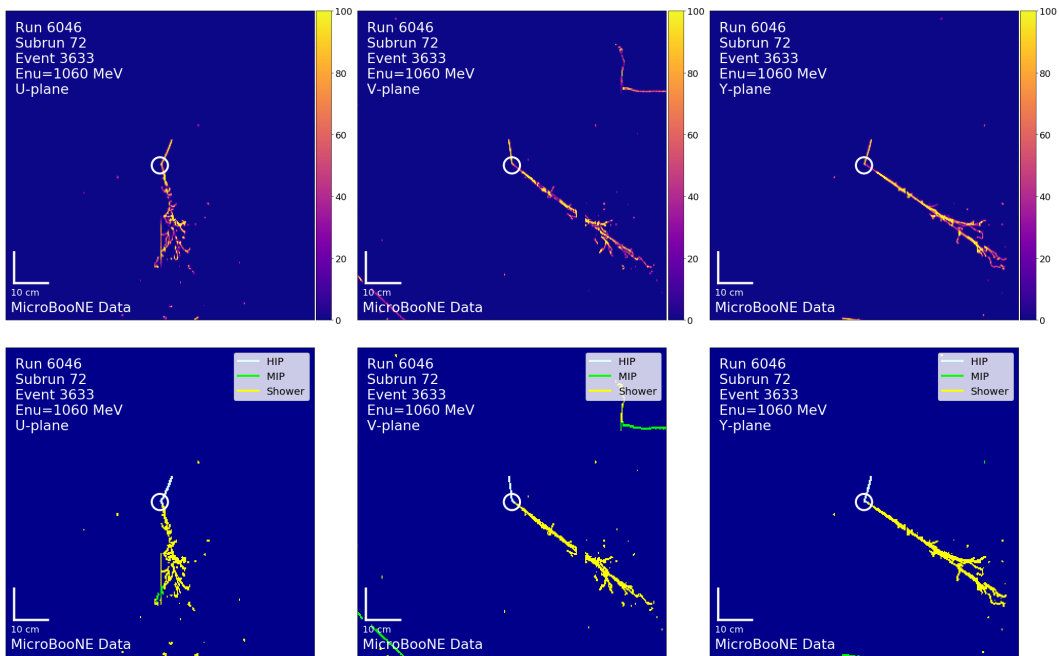


Figure B-24: Top: pixel intensity; Bottom: `SparseSSNet` labels; Left to right: U, V, Y, planes. The white circle indicates the reconstructed vertex.

Bibliography

- [1] John N. Bahcall. Solar models and solar neutrinos: Current status. *Phys. Scripta T*, 121:46–50, 2005.
- [2] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.
- [3] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002.
- [4] J. A. Formaggio and G. P. Zeller. From eV to EeV: Neutrino Cross Sections Across Energy Scales. *Rev. Mod. Phys.*, 84:1307–1341, 2012.
- [5] A. Aguilar-Arevalo et al. Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam. *Phys. Rev. D*, 64:112007, 2001.
- [6] K. N. Abazajian et al. Light Sterile Neutrinos: A White Paper. 4 2012.
- [7] C. Giunti, Y. F. Li, C. A. Ternes, O. Tyagi, and Z. Xin. Gallium Anomaly: critical view from the global picture of ν_e and $\bar{\nu}_e$ disappearance. *JHEP*, 10:164, 2022.
- [8] G. Mention, M. Fechner, Th. Lasserre, Th. A. Mueller, D. Lhuillier, M. Cri-bier, and A. Letourneau. The Reactor Antineutrino Anomaly. *Phys. Rev. D*, 83:073006, 2011.
- [9] V. V. Barinov et al. Results from the Baksan Experiment on Sterile Transitions (BEST). *Phys. Rev. Lett.*, 128(23):232501, 2022.
- [10] A. A. Aguilar-Arevalo et al. Updated MiniBooNE neutrino oscillation results with increased data and new background studies. *Phys. Rev. D*, 103(5):052002, 2021.
- [11] Mona Dentler, Álvaro Hernández-Cabezudo, Joachim Kopp, Pedro A. N. Machado, Michele Maltoni, Ivan Martinez-Soler, and Thomas Schwetz. Up-dated Global Analysis of Neutrino Oscillations in the Presence of eV-Scale Sterile Neutrinos. *JHEP*, 08:010, 2018.

- [12] J. M. Hardin, I. Martinez-Soler, A. Diaz, M. Jin, N. W. Kamp, C. A. Argüelles, J. M. Conrad, and M. H. Shaevitz. New Clues About Light Sterile Neutrinos: Preference for Models with Damping Effects in Global Fits. 11 2022.
- [13] A. A. Aguilar-Arevalo et al. Dark Matter Search in Nucleon, Pion, and Electron Channels from a Proton Beam Dump with MiniBooNE. *Phys. Rev. D*, 98(11):112004, 2018.
- [14] A. A. Aguilar-Arevalo et al. The Neutrino Flux prediction at MiniBooNE. *Phys. Rev. D*, 79:072002, 2009.
- [15] A. A. Aguilar-Arevalo et al. The MiniBooNE Detector. *Nucl. Instrum. Meth. A*, 599:28–46, 2009.
- [16] R. Acciarri et al. Design and Construction of the MicroBooNE Detector. *JINST*, 12(02):P02017, 2017.
- [17] Benjamin J. P. Jones. *Sterile Neutrinos in Cold Climates*. PhD thesis, MIT, 2015.
- [18] R. Acciarri et al. Noise Characterization and Filtering in the MicroBooNE Liquid Argon TPC. *JINST*, 12(08):P08003, 2017.
- [19] C. Adams et al. Ionization electron signal processing in single phase LArTPCs. Part I. Algorithm Description and quantitative evaluation with MicroBooNE simulation. *JINST*, 13(07):P07006, 2018.
- [20] None None. Microboone low-energy excess signal prediction from unfolding miniboone monte-carlo and data. 7 2018.
- [21] P. Abratenko et al. Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE. *Phys. Rev. D*, 103(5):052012, 2021.
- [22] P. Abratenko et al. Convolutional neural network for multiple particle identification in the MicroBooNE liquid argon time projection chamber. *Phys. Rev. D*, 103(9):092003, 2021.
- [23] Kiara Carloni, Nicholas W. Kamp, Austin Schneider, and Janet M. Conrad. Convolutional neural networks for shower energy prediction in liquid argon time projection chambers. *JINST*, 17(02):P02022, 2022.
- [24] P. Abratenko et al. Vertex-finding and reconstruction of contained two-track neutrino events in the MicroBooNE detector. *JINST*, 16(02):P02017, 2021.
- [25] Gary J. Feldman and Robert D. Cousins. A Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998.
- [26] S. S. Wilks. The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *Annals Math. Statist.*, 9(1):60–62, 1938.

- [27] Stefano Vergani, Nicholas W. Kamp, Alejandro Diaz, Carlos A. Argüelles, Janet M. Conrad, Michael H. Shaevitz, and Melissa A. Uchida. Explaining the MiniBooNE excess through a mixed model of neutrino oscillation and decay. *Phys. Rev. D*, 104(9):095005, 2021.
- [28] Lukas Zavorka, Michael J. Mocko, and Paul E. Koehler. Physics design of the next-generation spallation neutron target-moderator-reflector-shield assembly at lansce. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 901:189–197, 2018.
- [29] D. Baxter et al. Coherent Elastic Neutrino-Nucleus Scattering at the European Spallation Source. *JHEP*, 02:123, 2020.
- [30] A. A. Aguilar-Arevalo et al. First dark matter search results from Coherent CAPTAIN-Mills. *Phys. Rev. D*, 106(1):012001, 2022.
- [31] Nicholas W. Kamp, Matheus Hostert, Austin Schneider, Stefano Vergani, Carlos A. Argüelles, Janet M. Conrad, Michael H. Shaevitz, and Melissa A. Uchida. Dipole-coupled heavy-neutral-lepton explanations of the MiniBooNE excess including constraints from MINERvA data. *Phys. Rev. D*, 107(5):055009, 2023.
- [32] F. W. Clarke and H. S. Washington. The composition of the earth’s crust. Technical report, 1924.
- [33] W. Pauli. Dear radioactive ladies and gentlemen. *Phys. Today*, 31N9:27, 1978.
- [34] E. Fermi. An attempt of a theory of beta radiation. 1. *Z. Phys.*, 88:161–177, 1934.
- [35] F. Reines and C. L. Cowan. Detection of the free neutrino. *Phys. Rev.*, 92:830–831, 1953.
- [36] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956.
- [37] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Zh. Eksp. Teor. Fiz.*, 34:247, 1957.
- [38] G. Danby, J. M. Gaillard, Konstantin A. Goulianos, L. M. Lederman, Nari B. Mistry, M. Schwartz, and J. Steinberger. Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos. *Phys. Rev. Lett.*, 9:36–44, 1962.
- [39] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962.
- [40] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Zh. Eksp. Teor. Fiz.*, 53:1717–1725, 1967.

- [41] S. Bilenky. Neutrino oscillations: From a historical perspective to the present status. *Nucl. Phys. B*, 908:2–13, 2016.
- [42] Raymond Davis, Jr., Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, 1968.
- [43] John N. Bahcall, Neta A. Bahcall, and G. Shaviv. Present status of the theoretical predictions for the Cl-36 solar neutrino experiment. *Phys. Rev. Lett.*, 20:1209–1212, 1968.
- [44] D. Casper et al. Measurement of atmospheric neutrino composition with IMB-3. *Phys. Rev. Lett.*, 66:2561–2564, 1991.
- [45] K. S. Hirata et al. Observation of a small atmospheric muon-neutrino / electron-neutrino ratio in Kamiokande. *Phys. Lett. B*, 280:146–152, 1992.
- [46] John G. Learned. The Atmospheric neutrino anomaly: Muon neutrino disappearance. pages 89–130, 7 2000.
- [47] Y. Fukuda et al. The Super-Kamiokande detector. *Nucl. Instrum. Meth. A*, 501:418–462, 2003.
- [48] J. Boger et al. The Sudbury neutrino observatory. *Nucl. Instrum. Meth. A*, 449:172–207, 2000.
- [49] S. P. Mikheyev and A. Yu. Smirnov. Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos. *Sov. J. Nucl. Phys.*, 42:913–917, 1985.
- [50] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys. Rev. D*, 17:2369–2374, 1978.
- [51] T. Araki et al. Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion. *Phys. Rev. Lett.*, 94:081801, 2005.
- [52] Arthur B. McDonald. Nobel Lecture: The Sudbury Neutrino Observatory: Observation of flavor change for solar neutrinos. *Rev. Mod. Phys.*, 88(3):030502, 2016.
- [53] Takaaki Kajita. Nobel Lecture: Discovery of atmospheric neutrino oscillations. *Rev. Mod. Phys.*, 88(3):030501, 2016.
- [54] P. F. de Salas, D. V. Forero, S. Gariazzo, P. Martínez-Miravé, O. Mena, C. A. Ternes, M. Tórtola, and J. W. F. Valle. 2020 global reassessment of the neutrino oscillation picture. *JHEP*, 02:071, 2021.
- [55] Ivan Esteban, M. C. Gonzalez-Garcia, Michele Maltoni, Thomas Schwetz, and Albert Zhou. The fate of hints: updated global analysis of three-flavor neutrino oscillations. *JHEP*, 09:178, 2020.

- [56] Francesco Capozzi, Eleonora Di Valentino, Eligio Lisi, Antonio Marrone, Alessandro Melchiorri, and Antonio Palazzo. Unfinished fabric of the three neutrino paradigm. *Phys. Rev. D*, 104(8):083031, 2021.
- [57] K. Abe et al. Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations. *Nature*, 580(7803):339–344, 2020. [Erratum: *Nature* 583, E16 (2020)].
- [58] M. A. Acero et al. First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA. *Phys. Rev. Lett.*, 123(15):151803, 2019.
- [59] K. Abe et al. Physics potentials with the second Hyper-Kamiokande detector in Korea. *PTEP*, 2018(6):063C01, 2018.
- [60] R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF. 12 2015.
- [61] Fengpeng An et al. Neutrino Physics with JUNO. *J. Phys. G*, 43(3):030401, 2016.
- [62] R. N. Mohapatra and P. B. Pal. *Massive neutrinos in physics and astrophysics. Second edition*, volume 60. 1998.
- [63] C. H. Llewellyn Smith. Neutrino Reactions at Accelerator Energies. *Phys. Rept.*, 3:261–379, 1972.
- [64] Matthew D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, 3 2014.
- [65] Tsutomu Yanagida. Horizontal gauge symmetry and masses of neutrinos. *Conf. Proc. C*, 7902131:95–99, 1979.
- [66] M. Aker et al. Direct neutrino-mass measurement with sub-electronvolt sensitivity. *Nature Phys.*, 18(2):160–166, 2022.
- [67] S. Abe et al. Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen. *Phys. Rev. Lett.*, 130(5):051801, 2023.
- [68] C. Alduino et al. First Results from CUORE: A Search for Lepton Number Violation via $0\nu\beta\beta$ Decay of ^{130}Te . *Phys. Rev. Lett.*, 120(13):132501, 2018.
- [69] M. Agostini et al. Improved Limit on Neutrinoless Double- β Decay of ^{76}Ge from GERDA Phase II. *Phys. Rev. Lett.*, 120(13):132503, 2018.
- [70] Boris Kayser. On the Quantum Mechanics of Neutrino Oscillation. *Phys. Rev. D*, 24:110, 1981.
- [71] P. A. Zyla et al. Review of Particle Physics. *PTEP*, 2020(8):083C01, 2020.

- [72] G. S. Abrams et al. Measurements of Z Boson Resonance Parameters in e^+e^- Annihilation. *Phys. Rev. Lett.*, 63:2173, 1989.
- [73] M. Harada et al. Proposal: A Search for Sterile Neutrino at J-PARC Materials and Life Science Experimental Facility. 10 2013.
- [74] Takasumi Maruyama. The status of JSNS² and JSNS²-II. *PoS, NuFact2021*:159, 2022.
- [75] S. Ajimura et al. The JSNS2 detector. *Nucl. Instrum. Meth. A*, 1014:165742, 2021.
- [76] A. A. Aguilar-Arevalo et al. A Search for Electron Neutrino Appearance at the $\Delta m^2 \sim 1eV^2$ Scale. *Phys. Rev. Lett.*, 98:231801, 2007.
- [77] R. B. Patterson, E. M. Laird, Y. Liu, P. D. Meyers, I. Stancu, and H. A. Tanaka. The Extended-track reconstruction for MiniBooNE. *Nucl. Instrum. Meth. A*, 608:206–224, 2009.
- [78] A. A. Aguilar-Arevalo et al. Unexplained Excess of Electron-Like Events From a 1-GeV Neutrino Beam. *Phys. Rev. Lett.*, 102:101802, 2009.
- [79] A. A. Aguilar-Arevalo et al. Significant Excess of ElectronLike Events in the MiniBooNE Short-Baseline Neutrino Experiment. *Phys. Rev. Lett.*, 121(22):221801, 2018.
- [80] P. Abratenko et al. Search for Neutrino-Induced Neutral-Current Δ Radiative Decay in MicroBooNE and a First Test of the MiniBooNE Low Energy Excess under a Single-Photon Hypothesis. *Phys. Rev. Lett.*, 128:111801, 2022.
- [81] Vedran Brdar and Joachim Kopp. Can standard model and experimental uncertainties resolve the MiniBooNE anomaly? *Phys. Rev. D*, 105(11):115024, 2022.
- [82] Kevin J. Kelly and Joachim Kopp. More Ingredients for an Altarelli Cocktail at MiniBooNE. 10 2022.
- [83] G. Karagiorgi, B. R. Littlejohn, P. Machado, and Alexandre Sousa. Snowmass Neutrino Frontier: NF02 Topical Group Report on Understanding Experimental Neutrino Anomalies. In *2022 Snowmass Summer Study*, 9 2022.
- [84] Patrick Huber. On the determination of anti-neutrino spectra from nuclear reactors. *Phys. Rev. C*, 84:024617, 2011. [Erratum: *Phys.Rev.C* 85, 029901 (2012)].
- [85] Th. A. Mueller et al. Improved Predictions of Reactor Antineutrino Spectra. *Phys. Rev. C*, 83:054615, 2011.
- [86] M. Estienne et al. Updated Summation Model: An Improved Agreement with the Daya Bay Antineutrino Fluxes. *Phys. Rev. Lett.*, 123(2):022502, 2019.

- [87] L. Hayen, J. Kostensalo, N. Severijns, and J. Suhonen. First-forbidden transitions in the reactor anomaly. *Phys. Rev. C*, 100(5):054323, 2019.
- [88] V. Kopeikin, M. Skorokhvatov, and O. Titov. Reevaluating reactor antineutrino spectra with new measurements of the ratio between U235 and Pu239 β spectra. *Phys. Rev. D*, 104(7):L071301, 2021.
- [89] C. Giunti, Y. F. Li, C. A. Ternes, and Z. Xin. Reactor antineutrino anomaly in light of recent flux model refinements. *Phys. Lett. B*, 829:137054, 2022.
- [90] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate with gallium metal. III: Results for the 2002–2007 data-taking period. *Phys. Rev. C*, 80:015807, 2009.
- [91] F. Kaether, W. Hampel, G. Heusser, J. Kiko, and T. Kirsten. Reanalysis of the GALLEX solar neutrino flux and source experiments. *Phys. Lett. B*, 685:47–54, 2010.
- [92] Carlo Giunti and Marco Laveder. Short-Baseline Active-Sterile Neutrino Oscillations? *Mod. Phys. Lett. A*, 22:2499–2509, 2007.
- [93] A. Diaz, C. A. Argüelles, G. H. Collin, J. M. Conrad, and M. H. Shaevitz. Where Are We With Light Sterile Neutrinos? *Phys. Rept.*, 884:1–59, 2020.
- [94] Y. J. Ko et al. Sterile Neutrino Search at the NEOS Experiment. *Phys. Rev. Lett.*, 118(12):121802, 2017.
- [95] J. H. Choi et al. Search for Sub-eV Sterile Neutrinos at RENO. *Phys. Rev. Lett.*, 125(19):191801, 2020.
- [96] Feng Peng An et al. Improved Search for a Light Sterile Neutrino with the Full Configuration of the Daya Bay Experiment. *Phys. Rev. Lett.*, 117(15):151802, 2016.
- [97] H. Almazán et al. STEREO neutrino spectrum of ^{235}U fission rejects sterile neutrino hypothesis. *Nature*, 613(7943):257–261, 2023.
- [98] I. Alekseev et al. Optimized scintillation strip design for the DANSS upgrade. *JINST*, 17(04):P04009, 2022.
- [99] M. Andriamirado et al. Improved short-baseline neutrino oscillation search and energy spectrum measurement with the PROSPECT experiment at HFIR. *Phys. Rev. D*, 103(3):032001, 2021.
- [100] Z. Atif et al. Search for sterile neutrino oscillations using RENO and NEOS data. *Phys. Rev. D*, 105(11):L111101, 2022.
- [101] M. Aker et al. Improved eV-scale sterile-neutrino constraints from the second KATRIN measurement campaign. *Phys. Rev. D*, 105(7):072004, 2022.

- [102] A. A. Aguilar-Arevalo et al. MiniBooNE and MicroBooNE Combined Fit to a 3+1 Sterile Neutrino Scenario. *Phys. Rev. Lett.*, 129(20):201801, 2022.
- [103] P. Adamson et al. Search for sterile neutrinos in MINOS and MINOS+ using a two-detector fit. *Phys. Rev. Lett.*, 122(9):091803, 2019.
- [104] P. Adamson et al. Improved Constraints on Sterile Neutrino Mixing from Disappearance Searches in the MINOS, MINOS+, Daya Bay, and Bugey-3 Experiments. *Phys. Rev. Lett.*, 125(7):071801, 2020.
- [105] I. E. Stockdale et al. Limits on Muon Neutrino Oscillations in the Mass Range $55\text{-eV}^2 < \Delta m^2 < 800\text{-eV}^2$. *Phys. Rev. Lett.*, 52:1384, 1984.
- [106] M. G. Aartsen et al. eV-Scale Sterile Neutrino Search Using Eight Years of Atmospheric Muon Neutrino Data from the IceCube Neutrino Observatory. *Phys. Rev. Lett.*, 125(14):141801, 2020.
- [107] Janet M. Conrad, William C. Louis, and Michael H. Shaevitz. The LSND and MiniBooNE Oscillation Searches at High Δm^2 . *Ann. Rev. Nucl. Part. Sci.*, 63:45–67, 2013.
- [108] M. G. Catanesi et al. Measurement of the production cross-section of positive pions in the collision of 8.9-GeV/c protons on beryllium. *Eur. Phys. J. C*, 52:29–53, 2007.
- [109] I. Chemakin et al. Pion production by protons on a thin beryllium target at 6.4-GeV, 12.3-GeV/c, and 17.5-GeV/c incident proton momenta. *Phys. Rev. C*, 77:015209, 2008. [Erratum: *Phys.Rev.C* 77, 049903 (2008)].
- [110] C. L. Wang. Pion, kaon, and anti-proton production between 10 and 70 bev. *Phys. Rev. Lett.*, 25:1068–1072, 1970.
- [111] A. A. Aguilar-Arevalo et al. First Measurement of the Muon Neutrino Charged Current Quasielastic Double Differential Cross Section. *Phys. Rev. D*, 81:092005, 2010.
- [112] A. A. Aguilar-Arevalo et al. Measurement of muon neutrino quasi-elastic scattering on carbon. *Phys. Rev. Lett.*, 100:032301, 2008.
- [113] Alexis A. Aguilar-Arevalo et al. Measurement of ν_μ and $\bar{\nu}_\mu$ induced neutral current single π^0 production cross sections on mineral oil at $E_\nu \sim \mathcal{O}(1\text{GeV})$. *Phys. Rev. D*, 81:013005, 2010.
- [114] A. A. Aguilar-Arevalo et al. First Observation of Coherent π^0 Production in Neutrino Nucleus Interactions with $E_\nu < 2\text{ GeV}$. *Phys. Lett. B*, 664:41–46, 2008.
- [115] A. A. Aguilar-Arevalo et al. First measurement of the muon antineutrino double-differential charged-current quasielastic cross section. *Phys. Rev. D*, 88(3):032001, 2013.

- [116] P. Abratenko et al. Search for an Excess of Electron Neutrino Interactions in MicroBooNE Using Multiple Final-State Topologies. *Phys. Rev. Lett.*, 128(24):241801, 2022.
- [117] Yang Bai, Ran Lu, Sida Lu, Jordi Salvado, and Ben A. Stefanek. Three Twin Neutrinos: Evidence from LSND and MiniBooNE. *Phys. Rev. D*, 93(7):073004, 2016.
- [118] André de Gouvêa, O. L. G. Peres, Suprabh Prakash, and G. V. Stenico. On The Decaying-Sterile Neutrino Solution to the Electron (Anti)Neutrino Appearance Anomalies. *JHEP*, 07:141, 2020.
- [119] Mona Dentler, Ivan Esteban, Joachim Kopp, and Pedro Machado. Decaying Sterile Neutrinos and the Short Baseline Oscillation Anomalies. *Phys. Rev. D*, 101(11):115013, 2020.
- [120] Heinrich Pas, Sandip Pakvasa, and Thomas J. Weiler. Sterile-active neutrino oscillations and shortcuts in the extra dimension. *Phys. Rev. D*, 72:095017, 2005.
- [121] Marcela Carena, Ying-Ying Li, Camila S. Machado, Pedro A. N. Machado, and Carlos E. M. Wagner. Neutrinos in Large Extra Dimensions and Short-Baseline ν_e Appearance. *Phys. Rev. D*, 96(9):095014, 2017.
- [122] Dominik Döring and Heinrich Päs. Sterile Neutrino Shortcuts in Asymmetrically Warped Extra Dimensions. *Eur. Phys. J. C*, 79(7):604, 2019.
- [123] S. N. Gninenko. The MiniBooNE anomaly and heavy neutrino decay. *Phys. Rev. Lett.*, 103:241802, 2009.
- [124] Sergei N. Gninenko. A resolution of puzzles from the LSND, KARMEN, and MiniBooNE experiments. *Phys. Rev. D*, 83:015015, 2011.
- [125] Claudio Dib, Juan Carlos Helo, Sergey Kovalenko, and Ivan Schmidt. Sterile neutrino decay explanation of LSND and MiniBooNE anomalies. *Phys. Rev. D*, 84:071301, 2011.
- [126] S. N. Gninenko. New limits on radiative sterile neutrino decays from a search for single photons in neutrino interactions. *Phys. Lett. B*, 710:86–90, 2012.
- [127] Manuel Masip, Pere Masjuan, and Davide Meloni. Heavy neutrino decays at MiniBooNE. *JHEP*, 01:106, 2013.
- [128] Alexander Radionov. Constraints on electromagnetic properties of sterile neutrinos from MiniBooNE results. *Phys. Rev. D*, 88(1):015016, 2013.
- [129] Peter Ballett, Silvia Pascoli, and Mark Ross-Lonergan. MeV-scale sterile neutrino decays at the Fermilab Short-Baseline Neutrino program. *JHEP*, 04:102, 2017.

- [130] Gabriel Magill, Ryan Plestid, Maxim Pospelov, and Yu-Dai Tsai. Dipole Portal to Heavy Neutral Leptons. *Phys. Rev. D*, 98(11):115015, 2018.
- [131] A. Baha Balantekin, André de Gouvêa, and Boris Kayser. Addressing the Majorana vs. Dirac Question with Neutrino Decays. *Phys. Lett. B*, 789:488–495, 2019.
- [132] Shyam Balaji, Maura Ramirez-Quezada, and Ye-Ling Zhou. CP violation and circular polarisation in neutrino radiative decay. *JHEP*, 04:178, 2020.
- [133] Shyam Balaji, Maura Ramirez-Quezada, and Ye-Ling Zhou. CP violation in neutral lepton transition dipole moment. *JHEP*, 12:090, 2020.
- [134] Oliver Fischer, Álvaro Hernández-Cabezudo, and Thomas Schwetz. Explaining the MiniBooNE excess by a decaying sterile neutrino with mass in the 250 MeV range. *Phys. Rev. D*, 101(7):075045, 2020.
- [135] Luis Alvarez-Ruso and Eduardo Saul-Sala. Neutrino interactions with matter and the MiniBooNE anomaly. *Eur. Phys. J. ST*, 230(24):4373–4389, 2021.
- [136] Enrico Bertuzzo, Sudip Jana, Pedro A. N. Machado, and Renata Zukanovich Funchal. Neutrino Masses and Mixings Dynamically Generated by a Light Dark Sector. *Phys. Lett. B*, 791:210–214, 2019.
- [137] Peter Ballett, Silvia Pascoli, and Mark Ross-Lonergan. U(1)’ mediated decays of heavy sterile neutrinos in MiniBooNE. *Phys. Rev. D*, 99:071701, 2019.
- [138] Enrico Bertuzzo, Sudip Jana, Pedro A. N. Machado, and Renata Zukanovich Funchal. Dark Neutrino Portal to Explain MiniBooNE excess. *Phys. Rev. Lett.*, 121(24):241801, 2018.
- [139] Peter Ballett, Matheus Hostert, and Silvia Pascoli. Dark Neutrinos and a Three Portal Connection to the Standard Model. *Phys. Rev. D*, 101(11):115025, 2020.
- [140] Asli Abdullahi, Matheus Hostert, and Silvia Pascoli. A dark seesaw solution to low energy anomalies: MiniBooNE, the muon ($g-2$), and BaBar. *Phys. Lett. B*, 820:136531, 2021.
- [141] Alakabha Datta, Saeed Kamali, and Danny Marfatia. Dark sector origin of the KOTO and MiniBooNE anomalies. *Phys. Lett. B*, 807:135579, 2020.
- [142] Bhaskar Dutta, Sumit Ghosh, and Tianjun Li. Explaining $(g-2)_{\mu,e}$, the KOTO anomaly and the MiniBooNE excess in an extended Higgs model with sterile neutrinos. *Phys. Rev. D*, 102(5):055017, 2020.
- [143] Waleed Abdallah, Raj Gandhi, and Samiran Roy. Understanding the Mini-BooNE and the muon and electron $g2$ anomalies with a light Z and a second Higgs doublet. *JHEP*, 12:188, 2020.

- [144] Waleed Abdallah, Raj Gandhi, and Samiran Roy. Two-Higgs doublet solution to the LSND, MiniBooNE and muon $g-2$ anomalies. *Phys. Rev. D*, 104(5):055028, 2021.
- [145] Bhaskar Dutta, Doojin Kim, Adrian Thompson, Remington T. Thornton, and Richard G. Van de Water. Solutions to the MiniBooNE Anomaly from New Physics in Charged Meson Decays. *Phys. Rev. Lett.*, 129(11):111803, 2022.
- [146] Carlos A. Argüelles, Matheus Hostert, and Yu-Dai Tsai. Testing New Physics Explanations of the MiniBooNE Anomaly at Neutrino Scattering Experiments. *Phys. Rev. Lett.*, 123(26):261801, 2019.
- [147] Johnathon R. Jordan, Yonatan Kahn, Gordan Krnjaic, Matthew Moschella, and Joshua Spitz. Severe Constraints on New Physics Explanations of the MiniBooNE Excess. *Phys. Rev. Lett.*, 122(8):081801, 2019.
- [148] P. Abratenko et al. Search for an anomalous excess of charged-current quasielastic ν_e interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction. *Phys. Rev. D*, 105(11):112003, 2022.
- [149] W. J. Willis and V. Radeka. Liquid Argon Ionization Chambers as Total Absorption Detectors. *Nucl. Instrum. Meth.*, 120:221–236, 1974.
- [150] C. Rubbia. The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors. 5 1977.
- [151] D. R. Nygren. The Time Projection Chamber: A New 4 pi Detector for Charged Particles. *eConf*, C740805:58, 1974.
- [152] G. Charpak, D. Rahm, and H. Steiner. Some developments in the operation of multiwire proportional chambers. *Nucl. Instrum. Meth.*, 80:13–34, 1970.
- [153] S. Amerio et al. Design, construction and tests of the ICARUS T600 detector. *Nucl. Instrum. Meth. A*, 527:329–410, 2004.
- [154] C. Rubbia et al. Underground operation of the ICARUS T600 LAr-TPC: first results. *JINST*, 6:P07011, 2011.
- [155] M. Antonello et al. Search for anomalies in the ν_e appearance from a ν_μ beam. *Eur. Phys. J. C*, 73:2599, 2013.
- [156] R. Acciarri et al. Measurements of Inclusive Muon Neutrino and Antineutrino Charged Current Differential Cross Sections on Argon in the NuMI Antineutrino Beam. *Phys. Rev. D*, 89(11):112003, 2014.
- [157] B. Baibussinov et al. Free electron lifetime achievements in Liquid Argon Imaging TPC. *JINST*, 5:P03005, 2010.

- [158] B. J. P. Jones, C. S. Chiu, J. M. Conrad, C. M. Ignarra, T. Katori, and M. Toups. A Measurement of the Absorption of Liquid Argon Scintillation Light by Dissolved Nitrogen at the Part-Per-Million Level. *JINST*, 8:P07011, 2013. [Erratum: *JINST* 8, E09001 (2013)].
- [159] M. Suzuki and S. Kubota. MECHANISM OF PROPORTIONAL SCINTILLATION IN ARGON, KRYPTON AND XENON. *Nucl. Instrum. Meth.*, 164:197–199, 1979.
- [160] P. Adhikari et al. The liquid-argon scintillation pulse shape in DEAP-3600. *Eur. Phys. J. C*, 80(4):303, 2020.
- [161] S Kubota, M Hishida, and J Raun. Evidence for a triplet state of the self-trapped exciton states in liquid argon, krypton and xenon. *Journal of Physics C: Solid State Physics*, 11(12):2645, jun 1978.
- [162] Albert Einstein. Concerning an heuristic point of view toward the emission and transformation of light. *Annalen Phys.*, 17:132–148, 1905.
- [163] T. Brieser et al. Testing of Cryogenic Photomultiplier Tubes for the MicroBooNE Experiment. *JINST*, 8:T07005, 2013.
- [164] R. Acciarri et al. Oxygen contamination in liquid Argon: Combined effects on ionization electron charge and scintillation light. *JINST*, 5:P05003, 2010.
- [165] R. Acciarri et al. Effects of Nitrogen contamination in liquid Argon. *JINST*, 5:P06003, 2010.
- [166] J. Asaadi, B. J. P. Jones, A. Tripathi, I. Parmaksiz, H. Sullivan, and Z. G. R. Williams. Emanation and bulk fluorescence in liquid argon from tetraphenyl butadiene wavelength shifting coatings. *JINST*, 14(02):P02021, 2019.
- [167] Emily Grace and James A. Nikkel. Index of refraction, Rayleigh scattering length, and Sellmeier coefficients in solid and liquid argon and xenon. *Nucl. Instrum. Meth. A*, 867:204–208, 2017.
- [168] C. Adams et al. Ionization electron signal processing in single phase LArTPCs. Part II. Data/simulation comparison and performance in MicroBooNE. *JINST*, 13(07):P07007, 2018.
- [169] Norbert Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications*. The MIT Press, 08 1949.
- [170] P. Abratenko et al. Search for an anomalous excess of inclusive charged-current ν_e interactions in the MicroBooNE experiment using Wire-Cell reconstruction. *Phys. Rev. D*, 105(11):112005, 2022.
- [171] P. Abratenko et al. Search for an anomalous excess of charged-current ν_e interactions without pions in the final state with the MicroBooNE experiment. *Phys. Rev. D*, 105(11):112004, 2022.

- [172] P. Abratenko et al. Electromagnetic shower reconstruction and energy validation with Michel electrons and π^0 samples for the deep-learning-based analyses in MicroBooNE. *JINST*, 16(12):T12017, 2021.
- [173] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. Meth. A*, 614:87–104, 2010.
- [174] Costas Andreopoulos, Christopher Barry, Steve Dytman, Hugh Gallagher, Tomasz Golan, Robert Hatcher, Gabriel Perdue, and Julia Yarba. The GENIE Neutrino Monte Carlo Generator: Physics and User Manual. 10 2015.
- [175] Luis Alvarez-Ruso et al. Recent highlights from GENIE v3. *Eur. Phys. J. ST*, 230(24):4449–4467, 2021.
- [176] Júlia Tena-Vidal et al. Neutrino-nucleon cross-section model tuning in GENIE v3. *Phys. Rev. D*, 104(7):072009, 2021.
- [177] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. Inclusive Charged-Current Neutrino-Nucleus Reactions. *Phys. Rev. C*, 83:045501, 2011.
- [178] A. Bodek and J. L. Ritchie. Fermi Motion Effects in Deep Inelastic Lepton Scattering from Nuclear Targets. *Phys. Rev. D*, 23:1070, 1981.
- [179] V. R. Pandharipande and Steven C. Pieper. Nuclear transparency to intermediate-energy nucleons from (e, e'p) reactions. *Phys. Rev. C*, 45:791–798, 1992.
- [180] P. Abratenko et al. New $CC0\pi$ GENIE model tune for MicroBooNE. *Phys. Rev. D*, 105(7):072001, 2022.
- [181] Ko Abe et al. Measurement of double-differential muon neutrino charged-current interactions on C_8H_8 without pions in the final state using the T2K off-axis beam. *Phys. Rev. D*, 93(11):112012, 2016.
- [182] Artur M. Ankowski and Alexander Friedland. Assessing the accuracy of the GENIE event generator with electron-scattering data. *Phys. Rev. D*, 102(5):053001, 2020.
- [183] A. Papadopoulou et al. Inclusive Electron Scattering And The GENIE Neutrino Event Generator. *Phys. Rev. D*, 103:113003, 2021.
- [184] M. Khachatryan et al. Electron-beam energy reconstruction for neutrino oscillation measurements. *Nature*, 599(7886):565–570, 2021.
- [185] S. Agostinelli et al. GEANT4—a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250–303, 2003.
- [186] E.L. Snider and G. Petrillo. Larsoft: toolkit for simulation, reconstruction and analysis of liquid argon tpc neutrino detectors. *Journal of Physics: Conference Series*, 898(4):042057, oct 2017.

- [187] G. D’Agostini. A Multidimensional unfolding method based on Bayes’ theorem. *Nucl. Instrum. Meth. A*, 362:487–498, 1995.
- [188] C. A. Argüelles, I. Esteban, M. Hostert, Kevin J. Kelly, J. Kopp, P. A. N. Machado, I. Martinez-Soler, and Y. F. Perez-Gonzalez. MicroBooNE and the νe Interpretation of the MiniBooNE Low-Energy Excess. *Phys. Rev. Lett.*, 128(24):241802, 2022.
- [189] Omar Benhar, Alessandro Lovato, and Noemi Rocco. Contribution of two-particle–two-hole final states to the nuclear response. *Phys. Rev. C*, 92(2):024602, 2015.
- [190] Mark Thomson. *Modern particle physics*. Cambridge University Press, New York, 2013.
- [191] Artur M. Ankowski and Jan T. Sobczyk. Argon spectral function and neutrino interactions. *Phys. Rev. C*, 74:054316, 2006.
- [192] P. Abratenko et al. Neutrino event selection in the MicroBooNE liquid argon time projection chamber using Wire-Cell 3D imaging, clustering, and charge-light matching. *JINST*, 16(06):P06043, 2021.
- [193] P. Abratenko et al. Cosmic Ray Background Rejection with Wire-Cell LArTPC Event Reconstruction in the MicroBooNE Detector. *Phys. Rev. Applied*, 15(6):064071, 2021.
- [194] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [195] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [196] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9224–9232, 2018.
- [197] C. Adams et al. Deep neural network for pixel-level electromagnetic particle identification in the MicroBooNE liquid argon time projection chamber. *Phys. Rev. D*, 99(9):092001, 2019.
- [198] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.

- [199] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- [200] Yasser Ganjisaffar, Rich Caruana, and Cristina Videira Lopes. Bagging gradient-boosted trees for high precision, low variance ranking models. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '11*, page 8594, New York, NY, USA, 2011. Association for Computing Machinery.
- [201] Nicholas W. Kamp, Matheus Hostert, Carlos A. Argüelles, Janet M. Conrad, and Michael H. Shaevitz. Implications of MicroBooNE’s low sensitivity to electron antineutrino interactions in the search for the MiniBooNE excess. 1 2023.
- [202] J. E. Moyal. Stochastic processes and statistical physics. *Journal of the Royal Statistical Society: Series B (Methodological)*, 11(2):150–210, 1949.
- [203] Nicholas Higham. Cholesky factorization. *Wiley Interdisciplinary Reviews: Computational Statistics*, 1:251 – 254, 09 2009.
- [204] P. Abratenko et al. Novel approach for evaluating detector-related uncertainties in a LArTPC using MicroBooNE data. *Eur. Phys. J. C*, 82(5):454, 2022.
- [205] P. Abratenko et al. First Measurement of Inclusive Muon Neutrino Charged Current Differential Cross Sections on Argon at $E_\nu \sim 0.8$ GeV with the MicroBooNE Detector. *Phys. Rev. Lett.*, 123(13):131801, 2019.
- [206] Melanie Day and Kevin S. McFarland. Differences in Quasi-Elastic Cross-Sections of Muon and Electron Neutrinos. *Phys. Rev. D*, 86:053003, 2012.
- [207] J. Calcutt, C. Thorpe, K. Mahn, and Laura Fields. Geant4Reweight: a framework for evaluating and propagating hadronic interaction uncertainties in Geant4. *JINST*, 16(08):P08042, 2021.
- [208] Emanuel Parzen. On Estimation of a Probability Density Function and Mode. *The Annals of Mathematical Statistics*, 33(3):1065 – 1076, 1962.
- [209] Murray Rosenblatt. Remarks on Some Nonparametric Estimates of a Density Function. *The Annals of Mathematical Statistics*, 27(3):832 – 837, 1956.
- [210] Xiangpan Ji, Wenqiang Gu, Xin Qian, Hanyu Wei, and Chao Zhang. Combined Neyman–Pearson chi-square: An improved approximation to the Poisson-likelihood chi-square. *Nucl. Instrum. Meth. A*, 961:163677, 2020.
- [211] Thomas Junk. Confidence level computation for combining searches with small statistics. *Nucl. Instrum. Meth. A*, 434:435–443, 1999.
- [212] A L Read. Modified frequentist analysis of search results (the CL_s method). 2000.

- [213] Peter B. Denton. Sterile Neutrino Search with MicroBooNE’s Electron Neutrino Disappearance Data. *Phys. Rev. Lett.*, 129(6):061801, 2022.
- [214] P. Abratenko et al. First Constraints on Light Sterile Neutrino Oscillations from Combined Appearance and Disappearance Searches with the MicroBooNE Detector. *Phys. Rev. Lett.*, 130(1):011801, 2023.
- [215] Will Loinaz, Naotoshi Okamura, Saifuddin Rayyan, Tatsu Takeuchi, and L. C. R. Wijewardhana. The NuTeV anomaly, lepton universality, and nonuniversal neutrino gauge couplings. *Phys. Rev. D*, 70:113004, 2004.
- [216] T. Adams et al. Terascale Physics Opportunities at a High Statistics, High Energy Neutrino Scattering Experiment: NuSOng. *Int. J. Mod. Phys. A*, 24:671–717, 2009.
- [217] J. Park et al. Measurement of Neutrino Flux from Neutrino-Electron Elastic Scattering. *Phys. Rev. D*, 93(11):112007, 2016.
- [218] E. Valencia et al. Constraint of the MINER ν A medium energy neutrino flux using neutrino-electron elastic scattering. *Phys. Rev. D*, 100(9):092001, 2019.
- [219] L. Zazueta et al. Improved constraint on the MINER ν A medium energy neutrino flux using $\nu^- e^- \rightarrow \nu^- e^-$ data. *Phys. Rev. D*, 107(1):012001, 2023.
- [220] Robert E. Shrock. Electromagnetic Properties and Decays of Dirac and Majorana Neutrinos in a General Class of Gauge Theories. *Nucl. Phys. B*, 206:359–379, 1982.
- [221] Palash B. Pal and Lincoln Wolfenstein. Radiative Decays of Massive Neutrinos. *Phys. Rev. D*, 25:766, 1982.
- [222] K. S. Babu, Sudip Jana, and Manfred Lindner. Large Neutrino Magnetic Moments in the Light of Recent Experiments. *JHEP*, 10:040, 2020.
- [223] Howard Georgi and Michael E. Luke. Neutrino Moments, Masses and Custodial SU(2) Symmetry. *Nucl. Phys. B*, 347:1–11, 1990.
- [224] Manfred Lindner, Branimir Radovčić, and Johannes Welter. Revisiting Large Neutrino Magnetic Moments. *JHEP*, 07:139, 2017.
- [225] Carlo Giunti and Alexander Studenikin. Neutrino electromagnetic properties. *Phys. Atom. Nucl.*, 72:2089–2125, 2009.
- [226] Andre de Gouvea and James Jenkins. What can we learn from neutrino electron scattering? *Phys. Rev. D*, 74:033004, 2006.
- [227] A. B. Balantekin and N. Vassh. Magnetic moments of active and sterile neutrinos. *Phys. Rev. D*, 89(7):073013, 2014.

- [228] P. Vogel and J. Engel. Neutrino Electromagnetic Form-Factors. *Phys. Rev. D*, 39:3378, 1989.
- [229] Boris Kayser. Majorana Neutrinos and their Electromagnetic Properties. *Phys. Rev. D*, 26:1662, 1982.
- [230] K. S. Babu, Sudip Jana, Manfred Lindner, and Vishnu P. K. Muon $g - 2$ anomaly and neutrino magnetic moments. *JHEP*, 10:240, 2021.
- [231] S. N. Gninenko and N. V. Krasnikov. Limits on the magnetic moment of sterile neutrino and two photon neutrino decay. *Phys. Lett. B*, 450:165–172, 1999.
- [232] Pilar Coloma, Pedro A. N. Machado, Ivan Martinez-Soler, and Ian M. Shoemaker. Double-Cascade Events from New Physics in Icecube. *Phys. Rev. Lett.*, 119(20):201804, 2017.
- [233] Ryan Plestid. Luminous solar neutrinos I: Dipole portals. *Phys. Rev. D*, 104:075027, 2021.
- [234] Thomas Schwetz, Albert Zhou, and Jing-Yu Zhu. Constraining active-sterile neutrino transition magnetic moments at DUNE near and far detectors. *JHEP*, 21:200, 2020.
- [235] Mack Atkinson, Pilar Coloma, Ivan Martinez-Soler, Noemi Rocco, and Ian M. Shoemaker. Heavy Neutrino Searches through Double-Bang Events at Super-Kamiokande, DUNE, and Hyper-Kamiokande. *JHEP*, 04:174, 2022.
- [236] Patrick D. Bolton, Frank F. Deppisch, Kåre Fridell, Julia Harz, Chandan Hati, and Suchita Kulkarni. Probing active-sterile neutrino transition magnetic moments with photon emission from CE ν NS. *Phys. Rev. D*, 106(3):035036, 2022.
- [237] R. Andrew Gustafson, Ryan Plestid, and Ian M. Shoemaker. Neutrino Portals, Terrestrial Upscattering, and Atmospheric Neutrinos. 5 2022.
- [238] Maksym Ovchynnikov, Thomas Schwetz, and Jing-Yu Zhu. Dipole portal and neutrinophilic scalars at DUNE revisited: The importance of the high-energy neutrino tail. *Phys. Rev. D*, 107(5):055029, 2023.
- [239] Yu Zhang and Wei Liu. Probing active-sterile neutrino transition magnetic moments at LEP and CEPC. 1 2023.
- [240] Vedran Brdar, André de Gouvêa, Ying-Ying Li, and Pedro A. N. Machado. Neutrino magnetic moment portal and supernovae: New constraints and multimessenger opportunities. *Phys. Rev. D*, 107(7):073005, 2023.
- [241] R. N. Mohapatra. Mechanism for Understanding Small Neutrino Mass in Superstring Theories. *Phys. Rev. Lett.*, 56:561–563, 1986.
- [242] R. N. Mohapatra and J. W. F. Valle. Neutrino Mass and Baryon Number Nonconservation in Superstring Models. *Phys. Rev. D*, 34:1642, 1986.

- [243] Carlos A. Argüelles, Nicolò Foppiani, and Matheus Hostert. Heavy neutral leptons below the kaon mass at hodoscopic neutrino detectors. *Phys. Rev. D*, 105(9):095006, 2022.
- [244] K. Abe et al. Search for heavy neutrinos with the T2K near detector ND280. *Phys. Rev. D*, 100(5):052006, 2019.
- [245] A. V. Artamonov et al. Search for heavy neutrinos in $K^+ \rightarrow \mu^+ \nu_H$ decays. *Phys. Rev. D*, 91(5):052001, 2015. [Erratum: *Phys.Rev.D* 91, 059903 (2015)].
- [246] G. Bernardi et al. FURTHER LIMITS ON HEAVY NEUTRINO COUPLINGS. *Phys. Lett. B*, 203:332–334, 1988.
- [247] M. B. Voloshin. On Compatibility of Small Mass with Large Magnetic Moment of Neutrino. *Sov. J. Nucl. Phys.*, 48:512, 1988.
- [248] K. S. Babu and R. N. Mohapatra. Model for Large Transition Magnetic Moment of the ν_e . *Phys. Rev. Lett.*, 63:228, 1989.
- [249] Miriam Leurer and Neil Marcus. A Model for a Large Neutrino Magnetic Transition Moment and Naturally Small Mass. *Phys. Lett. B*, 237:81–87, 1990.
- [250] K. S. Babu and R. N. Mohapatra. Supersymmetry and Large Transition Magnetic Moment of the Neutrino. *Phys. Rev. Lett.*, 64:1705, 1990.
- [251] Riccardo Barbieri and Rabindra N. Mohapatra. A Neutrino With a Large Magnetic Moment and a Naturally Small Mass. *Phys. Lett. B*, 218:225–229, 1989.
- [252] S. Gariazzo, C. Giunti, M. Laveder, and Y. F. Li. Updated Global 3+1 Analysis of Short-BaseLine Neutrino Oscillations. *JHEP*, 06:135, 2017.
- [253] M. Maltoni and T. Schwetz. Testing the statistical compatibility of independent data sets. *Phys. Rev. D*, 68:033020, 2003.
- [254] David McKeen and Maxim Pospelov. Muon Capture Constraints on Sterile Neutrino Properties. *Phys. Rev. D*, 82:113018, 2010.
- [255] Ian M. Shoemaker and Jason Wyenberg. Direct Detection Experiments at the Neutrino Dipole Portal Frontier. *Phys. Rev. D*, 99(7):075010, 2019.
- [256] G. Fricke, C. Bernhardt, K. Heilig, L. A. Schaller, L. Schellenberg, E. B. Shera, and C. W. de Jager. Nuclear Ground State Charge Radii from Electromagnetic Interactions. *Atom. Data Nucl. Data Tabl.*, 60:177–285, 1995.
- [257] H. De Vries, C. W. De Jager, and C. De Vries. Nuclear charge and magnetization density distribution parameters from elastic electron scattering. *Atom. Data Nucl. Data Tabl.*, 36:495–536, 1987.

- [258] C. W. De Jager, H. De Vries, and C. De Vries. Nuclear charge and magnetization density distribution parameters from elastic electron scattering. *Atom. Data Nucl. Data Tabl.*, 14:479–508, 1974. [Erratum: *Atom. Data Nucl. Data Tabl.* 16, 580–580 (1975)].
- [259] <http://discovery.phys.virginia.edu/research/groups/ncd/index.html>.
- [260] R. Abbasi et al. LeptonInjector and LeptonWeighter: A neutrino event generator and weighter for neutrino observatories. *Comput. Phys. Commun.*, 266:108018, 2021.
- [261] CCM workshop 2023. <https://ccm.mit.edu/ccm-workshop-2023>.
- [262] LeptonInjector. <https://github.com/Harvard-Neutrino/LeptonInjector>.
- [263] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.
- [264] Asli M. Abdullahi, Jaime Hoefken Zink, Matheus Hostert, Daniele Massaro, and Silvia Pascoli. DarkNews: a Python-based event generator for heavy neutral lepton production in neutrino-nucleus scattering. 7 2022.
- [265] Z. Pavlovic, R.G. Van de Water, and S. Zeller. Miniboone gamma-ray and electron efficiencies. Technical report, 2012.
- [266] Private communication with MiniBooNE Collaboration.
- [267] Michael H. Shaevitz. MiniBooNE oscillation results and implications. *J. Phys. Conf. Ser.*, 120:052003, 2008.
- [268] MiniBooNE Collaboration. Updated MiniBooNE Neutrino Oscillation Results with Increased Data and New Background Studies. HEPData (collection), 2021. <https://doi.org/10.17182/hepdata.114365>.
- [269] D. Akimov et al. Observation of Coherent Elastic Neutrino-Nucleus Scattering. *Science*, 357(6356):1123–1126, 2017.
- [270] A. A. Aguilar-Arevalo et al. First Leptophobic Dark Matter Search from the Coherent-CAPTAIN-Mills Liquid Argon Detector. *Phys. Rev. Lett.*, 129(2):021801, 2022.
- [271] A. A. Aguilar-Arevalo et al. Axion-Like Particles at Coherent CAPTAIN-Mills. 12 2021.
- [272] A.T. Nelson, J.A. O’Toole, R.A. Valicenti, and S.A. Maloy. Fabrication of a tantalum-clad tungsten target for lances. *Journal of Nuclear Materials*, 431(1):172–184, 2012. Special Issue of the Tenth International Workshop on Spallation Materials Technology, (IWSMT-10).

- [273] Paul W. Lisowski and Kurt F. Schoenberg. The los alamos neutron science center. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 562(2):910–914, 2006. Proceedings of the 7th International Conference on Accelerator Applications.
- [274] H. Berns et al. The CAPTAIN Detector and Physics Program. In *Snowmass 2013: Snowmass on the Mississippi*, 9 2013.
- [275] Edward Dunton. *A Search for Axion-like Particles at the Coherent CAPTAIN Mills Experiment*. PhD thesis, Columbia University, 2022.
- [276] E. Simon et al. SICANE: A Detector array for the measurement of nuclear recoil quenching factors using a monoenergetic neutron beam. *Nucl. Instrum. Meth. A*, 507:643–656, 2003.
- [277] D. Akimov et al. First Measurement of Coherent Elastic Neutrino-Nucleus Scattering on Argon. *Phys. Rev. Lett.*, 126(1):012002, 2021.
- [278] C. E. Taylor et al. The Mini-CAPTAIN liquid argon time projection chamber. *Nucl. Instrum. Meth. A*, 1001:165131, 2021.
- [279] Claus Grupen and Boris Schwartz. *Particle detectors*. Cambridge Univ. Pr., Cambridge, UK, 2008.
- [280] T. Kaptanoglu, E. J. Callaghan, M. Yeh, and G. D. Orebi Gann. Cherenkov and scintillation separation in water-based liquid scintillator using an LAPPDTM. *Eur. Phys. J. C*, 82(2):169, 2022.
- [281] J. Caravaca, B. J. Land, M. Yeh, and G. D. Orebi Gann. Characterization of water-based liquid scintillator for Cherenkov and scintillation separation. *Eur. Phys. J. C*, 80(9):867, 2020.
- [282] Jack Dunger, Edward J. Leming, and Steven D. Biller. Slow-fluor scintillator for low energy solar neutrinos and neutrinoless double beta decay. *Phys. Rev. D*, 105(9):092006, 2022.
- [283] Julieta Gruszko, Brian Naranjo, Byron Daniel, Andrey Elagin, Diana Gooding, Chris Grant, Jonathan Ouellet, and Lindley Winslow. Detecting Cherenkov light from 1–2 MeV electrons in linear alkylbenzene. *JINST*, 14(02):P02005, 2019.
- [284] M. Agostini et al. First Directional Measurement of Sub-MeV Solar Neutrinos with Borexino. *Phys. Rev. Lett.*, 128(9):091803, 2022.
- [285] M. Antonello et al. Detection of Cherenkov light emission in liquid argon. *Nucl. Instrum. Meth. A*, 516:348–363, 2004.

- [286] Christopher Benson, Gabriel Orebi Gann, and Victor Gehman. Measurements of the intrinsic quantum efficiency and absorption length of tetraphenyl butadiene thin films in the vacuum ultraviolet regime. *Eur. Phys. J. C*, 78(4):329, 2018.
- [287] Junghyun Bae and Stylianos Chatzidakis. Fieldable muon spectrometer using multi-layer pressurized gas Cherenkov radiators and its applications. *Sci. Rep.*, 12(1):2559, 2022.
- [288] J.M. Flourney, I.B. Berlman, B. Rickborn, and R. Harrison. Substituted tetraphenylbutadienes as fast scintillator solutes. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 351(2):349–358, 1994.
- [289] Hamamatsu Photonics. *Photomultiplier tubes*. Hamamatsu, 2000.
- [290] M. G. Aartsen et al. In-situ calibration of the single-photoelectron charge response of the IceCube photomultiplier tubes. *JINST*, 15(06):P06032, 2020.
- [291] Cosmicwatch. <http://www.cosmicwatch.lns.mit.edu/>.