Abstract

First Measurement of Inclusive Muon Neutrino Charged Current Triple-Differential Cross Section on Argon

London James Cooper-Troendle

2023

The field of accelerator neutrino experiments is entering an era of precision oscillation measurements in which many of the remaining unknown neutrino measurements will be determined. The upcoming DUNE and Hyper-K experiments aim to determine the neutrino mass hierarchy and degree of Charge-Parity (CP) violation in the neutrino sector, providing potential insight on the matter-antimatter imbalance observed in the universe. However, these experiments require highly accurate measurements, and neutrino cross section modeling uncertainties may limit their capabilities. Cross section measurements at current-generation experiments can aid the development of neutrino interaction models to reduce these uncertainties. This is especially true for measurements of neutrino energy, which drive neutrino oscillations and are of key importance to oscillation experiments.

The MicroBooNE experiment uses a Liquid Argon Time Projection Chamber (LArTPC) to produce neutrino-argon cross sections as one of its physics goals. The MicroBooNE detector's fully active volume, precision reconstruction, and calorimetry information are leveraged in the Wire-Cell analysis to produce a muon neutrino selection that is 92% pure while maintaining 68% efficiency. A reconstruction chain featuring a fully 3D charge reconstruction and a graph-based particle trajectory fit are used to produce accurate measurements of lepton kinematics as well as visible hadronic energy produced in a neutrino interaction. This thesis presents the first neutrino-argon triple-differential cross section measurement, targeting inclusive charged-current final states. Wiener SVD unfolding is used to produce a measurement over neutrino energy, muon momentum, and muon scattering angle. A series of constrained goodness of fit tests are used to demonstrate the validity of MicroBooNE's model in describing

the distribution of reconstructed kinematics seen in data to ensure the accuracy of unfolding. The validated unfolding to neutrino energy represents a step forward in the field of neutrino cross sections, and demonstrates the capabilities of the LArTPC detector.

First Measurement of Inclusive Muon Neutrino Charged Current Triple-Differential Cross Section on Argon

A Dissertation

Presented to the Faculty of the Graduate School

of

Yale University

In Candidacy for the Degree of

Doctor of Philosophy

By

London Cooper-Troendle

Dissertation Directors: Bonnie Fleming and Karsten Heeger

May 2023

Contents

1 Theory **1**
1.1 Neutrino Mass . . . 2
1.2 Neutrino Mixing . . . 6
1.3 Neutrino Oscillations . . . 7
1.4 Two-Flavor Oscillations . . . 10

2 Neutrino Interactions **12**
2.1 Importance of Cross Section Measurements . . . 12
2.2 Neutrino Interaction Channels . . . 13
2.3 Nuclear Effects . . . 17

3 History of Neutrino Experiments **20**
3.1 Discovery . . . 20
3.2 Oscillation Parameter Searches . . . 22
3.3 Overview of Detector Types . . . 27

4 The MicroBooNE Experiment **30**
4.1 Primary Goals . . . 30
4.2 The Booster Neutrino Beam . . . 32
4.3 The MicroBooNE LArTPC . . . 34
4.3.1 Liquid Argon Physics . . . 35
4.3.2 Recombination . . . 36
4.3.3 Diffusion . . . 38
4.3.4 The Space Charge Effect . . . 39
4.3.5 Charge Measurement . . . 40
4.3.6 Light Collection System . . . 43

5 Wire-Cell Event Reconstruction **44**
5.1 Signal Processing . . . 44
5.1.1 Field Response . . . 45
5.1.2 Charge Measurement . . . 50
5.2 Tomographic Imaging . . . 54
5.2.1 Tiling . . . 55
5.2.2 Charge Solving . . . 58
5.2.3 Deghosting . . . 62
5.2.4 Clustering . . . 63
5.3 Charge-Light Matching . . . 69
5.3.1 Flash Reconstruction . . . 69
5.3.2 Match Determination . . . 70
5.4 Trajectory Fitting . . . 75
5.4.1 Minimum Spanning Tree . . . 77
5.4.2 Trajectory Seed . . . 78
5.4.3 Trajectory Fit . . . 79

5.4.4 dQ/dx Fit . . . 80
5.5 Particle Identification and Vertexing . . . 83
5.5.1 Track Segmentation . . . 83
5.5.2 Vertex Fitting . . . 84
5.5.3 Track-Shower Separation . . . 85
5.5.4 Traditional Neutrino Vertex Identification . . . 86
5.5.5 Deep Neutral Network Vertexing . . . 88
5.5.6 Particle Flow Tree . . . 91
5.5.7 Energy Reconstruction . . . 94

6 Wire-Cell Event Selection **96**
6.1 Generic Neutrino Detection . . . 96
6.2 ν_μCC Selection . . . 102
6.3 Event Selection Over 1D Distributions . . . 102
6.4 Event Selection Over 3D Distribution . . . 108

7 Resolutions of Reconstructed Quantities **121**
7.1 Performance of Charge Reconstruction . . . 121
7.2 Vertex Identification Resolution . . . 124
7.3 Resolutions of Kinematic Variables . . . 128

8 Estimation of Systematic Uncertainties **130**
8.1 Estimation of Monte-Carlo Statistical Uncertainties . . . 130
8.2 Estimation of Flux Uncertainties . . . 131
8.3 Estimation of Cross Section Uncertainties . . . 133
8.4 Estimation of Detector Systematic Uncertainties . . . 136
8.4.1 Bootstrapping . . . 137
8.4.2 Gaussian Processes Regression Smoothing . . . 139
8.5 Total Covariance Matrix . . . 142

9 Model Validation **145**
9.1 The Wire-Cell Model Validation Procedure . . . 148
9.2 Validation of the Modeling of Kinematic Event Distributions . . . 152
9.3 Validation of the Modeling of Partially Contained Event Reconstruction . . . 154
9.4 Validation of the Modeling of Transfer Energy . . . 156
9.5 Examining the Goodness of Fit Test with Fake Data . . . 159

10 Unfolding Procedure **164**
10.1 Unfolding without Regularization . . . 164
10.2 Wiener SVD Unfolding . . . 166
10.3 Regularization in Multiple Dimensions . . . 169
10.4 Estimation of the Detector Response Matrix . . . 170

11 Cross Section Measurements **174**
11.1 Single-Differential Cross Section Measurements 174
11.2 Triple-Differential Cross Section Measurement 175
11.3 Conclusion 179

Acknowledgements

I would like to thank my advisor Bonnie Fleming for her continuing support throughout my PhD. She initially inspired me to work in the field of neutrino physics, and continued to guide and encourage me over the last six years. She has been a role model, all while making sure my career has been on a strong path.

I would also like to thank Xin Qian, who has led the Wire Cell group and served as an unofficial second advisor to me. He has guided a visionary new approach to LArTPC data reconstruction and motivated me to help contribute to the project. By striving for the highest standard of results, he has ensured that my work has been exciting, rewarding, and something I am proud of.

Additionally, I would like to thank Karsten Heeger for stepping in as formal advisor after Bonnie accepted the role of Fermilab deputy director. Karsten has helped me coordinate my defense and ensure that my path to graduation proceeds smoothly.

I'm grateful for all of the help I've received from close coworkers and friends over the years. Thanks Xiao for helping to introduce me to MicroBooNE and my first analysis project, when I probably took more work to teach than I could contribute. Brooke, thanks for always taking the time to explain the complicated details in Wire-Cell, even as you were trying to graduate, yourself. Giacomo, thanks for being a great friend to live with and also to help each other out on our projects through our PhDs. Thank you to Jay, Hanyu, Xiangpan, and Nitish for making so many steps of this analysis possible, and explaining to me how they work. And especially, thanks for Wenqiang for always volunteering to help produce a plot, debug a program, or explain a topic to me, even at odd hours of the day.

Finally, thank you mom and dad for the immediate support over the course of my PhD, as well as over my entire childhood to help me become a scientist. You instilled the sense of curiosity in me, and supported my growth that has led me here; I wouldn't have been able to do any of it without you.

Forward

My dissertation focuses on the steps involved in producing the first neutrino-argon inclusive muon-neutrino charged-current triple-differential cross section measurement. This measurement can help improve neutrino interaction modeling, assisting future neutrino oscillation experiments such as DUNE. Sec. 1 provides a theoretical framework to neutrino physics, and Sec. 2 discusses the importance and modeling of neutrino interactions. Sec. 3 gives an overview of the history of neutrino experiments, which gives context to the MicroBooNE experiment, discussed in general detail in Sec. 4. This analysis focuses on the Wire-Cell reconstruction and selection within MicroBooNE, which are extensively described in Sec. 5 and Sec. 6, respectively. The measurement resolutions of various reconstructed quantities are discussed in Sec. 7, and the estimation of various systematic uncertainties on the MicroBooNE model prediction are covered in Sec. 8. The validation of this model is introduced in Sec. 9, based on a new data-driven conditional constraint procedure. With the model fully introduced and validated, the reconstructed event selection is unfolded, following the Wiener SVD technique discussed in Sec. 10, to produce the triple-differential cross section measurement presented in Sec. 11.

My personal involvement includes contributions to the reconstruction, estimation of uncertainties, and unfolding. I worked on improving the flash reconstruction algorithm discussed in Sec. 5.3.1 as well as on the light mismatch algorithm designed to help identify and remove cosmic ray events, discussed in Sec. 6.1. I helped improve the estimation of systematic uncertainties through the development of the smoothing algorithm used to reduce statistical fluctuations involved in bootstrapping, discussed in Sec. 8.4.2. Finally, my largest contributions are specific to the triple-differential analysis, including determining the choice of binning in Sec. 6.4, validating the model prediction in Sec. 9.2-9.4, writing some of the unfolding code aimed at multi-dimensional unfolding, especially Sec. 10.3 , and producing the triple-differential measurement in Sec. 11.2.

1 Theory

The standard model of particle physics is one of the must successful theories in physics for predicting a wide range of experimental results. It describes the strong, weak, and electromagmetic forces, as well as all 17 known fundamental particles, including the six quarks, three leptons and their respective neutrinos, four force carriers, and the Higgs boson. Neutrinos were first hypothesized almost 100 years ago at a relatively early time within standard model particle physics, which was not fully developed until the discovery of quarks in the late 20th century. However, today neutrinos remain some of the least understood particles, with some of their properties not yet determined. Over the last century studying neutrinos has repeatedly revealed insight into new physics, and current and upcoming neutrino experiments hope to provide further insight through measuring the remaining unknowns.

In 1930, Wolfgang Pauli proposed the existence of the neutrino as a way to preserve conservation of energy and momentum in beta decay processes, where a neutron decays to a proton and emits an electron as well as an anti-electron-neutrino:

$$n \rightarrow p + e^- + \overline{\nu}_e \tag{1.1}$$

Experiments showed a continuous distribution of ejected electron energies, suggesting that the total energy released on decay was split among the electron and another particle, named the neutron by Pauli (and later renamed the neutrino). Furthermore, although electrons carry spin $\frac{1}{2}$, beta decay creates no net change in the spin of the nucleon, so the production of a neutrino with opposite spin to the electron is necessary to conserve angular momentum. The neutrino's existence was experimentally confirmed in 1956 by Cowan and Reines [1] by observing the capture of reactor-produced anti-neutrinos on protons:

$$p + \overline{\nu}_e \rightarrow n + e^+ \tag{1.2}$$

To account for the vanishingly small neutrino interaction probabilities, a new interaction mechanism was required, dubbed the weak force. This force is mediated by charged massive bosons W^+ and W^- as well as a neutral massive boson Z^0, and mediates interactions between all of the quarks and leptons, not just neutrinos [2]. Only left-handed neutrinos have ever been observed [3], a fact which initially motivated the massless description of the neutrino in the standard model [4]. An intuitive interpretation is that if neutrinos travel below the speed of light (and therefore have mass), it would be possible to change frame of reference to flip their helicity; only a massless neutrino can exist with only a single handedness. However, the standard model massless description of neutrinos and the conservation of lepton flavor is only an accidental symmetry, not one imposed by the model, and so it is possible for new physics to break the symmetry through a more detailed description of the physics.

1.1 Neutrino Mass

With the discovery of neutrino oscillations the existence of neutrino mass is established, and the presumed framework used in the standard model must be expanded [5]. Neutrino oscillations depend on differences in mass between neutrino mass eigenstates, so the measurement of oscillations implies that at most one neutrino state can be mass-less, and in all three may have mass. There are multiple mechanisms that could be used to explain the existence of neutrino mass [6, 7]. Some of them are less compelling, such as the prospect that the familiar Higgs interaction gives the neutrino mass like it does to other standard model particles, but for an unexplained reason with a much weaker coupling. Perhaps the most plausible explanation is through the Seesaw Mechanism [8, 9, 10], which explains neutrino mass without introducing any new fields. In this explanation, the Higgs boson couples only to the right-handed neutrino of each flavor. Then, slight mixing of the mass and flavor eigenstates allows for a massive but sterile (non-weakly-interacting) and thus unobserved neutrino, as well as the lightweight left-handed neutrino that has been experimentally detected.

The right- and left-handed wavefunctions Ψ_R and Ψ_L are eigenstates of the chirality matrix $\gamma^5 \equiv i\gamma^0\gamma^1\gamma^2\gamma^3$, where γ^0 through γ^3 are the gamma matrices:

$$\gamma^5 \Psi_{R/L} = \pm \Psi_{R/L} \tag{1.3}$$

This enables the construction of chiral projection operators P_R and P_L, which project the wavefunction into its right- and left-handed states:

$$P_{R/L} = \frac{I \pm \gamma^5}{2} \tag{1.4}$$

$$P_{R/L}\Psi = \Psi_{R/L} \tag{1.5}$$

$$\Psi = P_R\Psi + P_L\Psi = \Psi_R + \Psi_L \tag{1.6}$$

Antiparticles are described by the conjugation operator $\hat{C} : \Psi \rightarrow \Psi^c = C\overline{\Psi}^T$ with the conjugation matrix $C = i\gamma^2\gamma^0$. This produces the expected result that right- and left-handed fermions are antiparticles of each other:

$$(\Psi)^c_{R/L} = (\Psi^c)_{L/R} \tag{1.7}$$

By substituting the right- and left-handed fields into the Dirac Lagrangian, the relationship between Dirac mass term and the right- and left-handed fields is found:

$$\mathscr{L}_D = \overline{\Psi}_R i\gamma^\mu \partial_\mu \overline{\Psi}_R + \overline{\Psi}_L i\gamma^\mu \partial_\mu \overline{\Psi}_L - m_D(\overline{\Psi}_R\Psi_L + \overline{\Psi}_L\Psi_R) \tag{1.8}$$

Furthermore, the electroweak Lagrangian also allows for right and left Majorana mass terms

m_R and m_L (not to be confused with Majorana and Dirac particle types):

$$\mathscr{L}_M = -\frac{1}{2} m_R (\overline{\Psi}_R^c \Psi_R + \overline{\Psi}_R \Psi_R^c) - \frac{1}{2} m_L (\overline{\Psi}_L^c \Psi_L + \overline{\Psi}_L \Psi_L^c) \tag{1.9}$$

The distinction between Majorana and Dirac mass terms in the Lagrangian is useful for looking at how they interact with weak charge and leptonic (non) conservation. Assuming that the neutrino is a Dirac particle, meaning that ν and $\overline{\nu}$ are separate particles, the Dirac Lagrangian term violates weak charge by converting ν_L with weak charge $\frac{1}{2}$ to the sterile ν_R with weak charge 0 (or vice versa). However, since this process does not, for example, replace ν with $\overline{\nu}$, lepton number is conserved. This is contrasted with the Majorana terms, which allow for the creation of neutrinos from the vacuum, violating weak charge and lepton number. Things change slightly if the neutrino is a Majorana particle, where ν and $\overline{\nu}$ are the same particle and therefore must have weak charge of 0, nullifying the weak charge violation present in the change of handedness. In either case, the potential Dirac and Majorana mass terms can be grouped to form the matrix depiction:

$$\mathscr{L}_{mass} = -\frac{1}{2} [\overline{\Psi}_L, \overline{\Psi}_R^c] M \begin{bmatrix} \Psi_L^c \\ \Psi_R \end{bmatrix} - \frac{1}{2} [\overline{\Psi}_L^c, \overline{\Psi}_R] M \begin{bmatrix} \Psi_L \\ \Psi_R^c \end{bmatrix} \tag{1.10}$$

$$M = \begin{bmatrix} m_L & m_D \\ m_D & m_R \end{bmatrix} \tag{1.11}$$

So far neutrinos have been discussed in their flavor states Ψ_R and Ψ_L. Switching to mass states, the Seesaw Mechanism supposes that the Higgs boson only couples to one of the neutrino mass eigenstates Ψ_α and Ψ_β, generating eigenvalues of 0 and a GUT-scale [11] mass term m_ν, respectively. Solving Eqn 1.11 for the mass eigenvalues 0 and m_ν in terms of m_R, m_L, and m_D gives:

$$m_D^2 = m_R m_L \tag{1.12}$$

$$m_\nu = m_R + m_L \tag{1.13}$$

Equation 1.12 is what gives rise to the name of the Seesaw Mechanism, as the larger m_R gets, the smaller m_D must be to counterbalance it. Furthermore, solving for the neutrino mass eigenstates, we see that they are composed of only slightly mixed combinations of the flavor states:

$$\Psi_\alpha = (\Psi_L + \Psi_L^c) - \frac{m_D}{m_R}(\Psi_R + \Psi_R^c) \tag{1.14}$$

$$\Psi_\beta = (\Psi_R + \Psi_R^c) + \frac{m_D}{m_R}(\Psi_L + \Psi_L^c) \tag{1.15}$$

Given the incredibly small allowed values for m_L and the imagined massive GUT-scale m_R, we necessarily find that $m_D << m_R$, giving rise to a small degree of mixing. Ultimately, the Seesaw Mechanism proposes that this small mass eigenstate mixing gives rise to the small left-handed neutrino mass that is observed.

Unfortunately, the Seesaw Mechanism can only be tested at much higher energies, placing it outside the reach of current experimental capabilities. However, in the event that there are Majorana neutrinos [12], their mass can be tested through neutrinoless double beta decay [13]:

$$(Z, N) \rightarrow (Z+2, N-2) + 2e^- \tag{1.16}$$

Under double beta decay, two neutrons decay to protons in a nucleus, emitting electrons as well as anti-electron neutrinos. However, in the hypothetical case of Majorana type neutrinos, they are their own anti-particles and are capable of annihilating each other, leaving no neutrino emission. This process inherently violates lepton number, and even lepton minus baryon number L-B. To conserve angular momentum, the neutrinos can only annihilate if one is left-handed and the other right-handed, which has a probability of occurrence proportional to m_ν / E. Therefore, experiments measuring this rate are sensitive to the overall

neutrino mass. However, it is a second order effect due to the simultaneous decay of two neutrons, making it a rare process to observe. Furthermore, the low energies involved make neutrino detection (or rather, confidence of their absence) difficult. Still, experiments have set a lower bound on the decay lifetime at around 10^{26} seconds, and experiments such as CUORE [14] and KamLAND-Zen [15] continue to search for the existence of the Majorana neutrino through neutrinoless double beta decay.

1.2 Neutrino Mixing

Although the mechanism giving neutrinos mass is not yet determined, the oscillations generated by neutrino mass differences are well understood. There are three neutrino flavors $\nu_\alpha \in \{\nu_e, \nu_\mu, \nu_\tau\}$, corresponding to the three lepton flavors $\alpha \in \{e, \mu, \tau\}$. Similar to the quark sector, neutrino flavors are a mixture of their mass eignestates $\nu_j \in \{\nu_1, \nu_2, \nu_3\}$. While quark mixing is described through the CKM matrix [16, 17], neutrino mixing is described through the PMNS matrix U, named after Bruno Pontecorvo for predicting neutrino oscillations [18] and Ziro Maki, Masami Nakagawa, and Shoichi Sakata for constructing the mixing matrix formalism [19]. This construction allows any flavor state to be expressed as a linear combination of mass eigenstates, and vice-versa:

$$|\nu_\alpha\rangle = \sum_j U_{\alpha j} \left|\nu_j\right\rangle \tag{1.17}$$

$$\left|\nu_j\right\rangle = \sum_\alpha U^*_{\alpha j} |\nu_\alpha\rangle \tag{1.18}$$

For three neutrino flavors, the PMNS matrix is a 3×3 unitary matrix described by nine degrees of freedom. However, five degrees of freedom can be absorbed into the phases of particle fields, leaving a degree of freedom for the overall phase. As a result there are a total of four degrees of freedom [20], which can be described through three mixing angles,

$\{\theta_{12}, \theta_{13}, \theta_{23}\}$, and a complex phase δ_{CP} that allows for potential Charge-Parity (CP) violation:

$$
\begin{aligned}
U &= \begin{bmatrix} U_{11} & U_{12} & U_{13} \\ U_{21} & U_{22} & U_{23} \\ U_{31} & U_{32} & U_{33} \end{bmatrix} \\
&= \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix} \\
&= \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{bmatrix}
\end{aligned}
\tag{1.19}
$$

Here s_{ij} and c_{ij} represent $\sin\theta_{ij}$ and $\cos\theta_{ij}$ respectively.

1.3 Neutrino Oscillations

Although neutrinos are always created and detected in flavor states, their time evolution is governed by their energy, which depends on their mass eigenstates. As a result, neutrino mass differences drive oscillations between flavor states as the phases evolve at different rates. In all known cases neutrinos are ultra-relativistic, with their momentum p_j much larger than their mass m_j. This allows for the energy E_j of each mass eigenstate to be approximated as:

$$
E_j = \sqrt{p_j^2 + m_j^2} \approx p_j + \frac{m_j^2}{2p_j} \approx E + \frac{m_j^2}{2E} \tag{1.20}
$$

where the neutrino energy E was substituted for the momentum p_j. The relativistic substitution of $t \approx L$ can help replace the unobserved time t in the neutrino's frame with the measurable distance travelled L. In the case of vacuum oscillations, the time evolution of the

mass eignestates becomes:

$$\left|\nu_j(t)\right\rangle = e^{-iHt}\left|\nu_j(0)\right\rangle = e^{-i(E+\frac{m_j^2}{2E})L}\left|\nu_j(0)\right\rangle \approx e^{-i\frac{m_j^2 L}{2E}}\left|\nu_j(0)\right\rangle \tag{1.21}$$

where the last step ignores the overall common phase e^{-iEL} of the eigenstates. From this, it is possible to work out the oscillation probability from an initial flavor state α to a final state β:

$$\begin{aligned}
P_{\nu_\alpha\to\nu_\beta}(L) &= \left|\left\langle\nu_\beta\middle|\nu_\alpha(L)\right\rangle\right|^2 \\
&= \left|\left\langle\nu_\beta\middle|\sum_j U_{\alpha j}\middle|\nu_j(L)\right\rangle\right|^2 \\
&= \left|\left\langle\nu_\beta\middle|\sum_j U_{\alpha j}e^{-i\frac{m_j^2 L}{2E}}\middle|\nu_j(0)\right\rangle\right|^2 \\
&= \left|\left\langle\nu_\beta\middle|\sum_j\sum_\gamma U^*_{\gamma j}U_{\alpha j}e^{-i\frac{m_j^2 L}{2E}}\middle|\nu_\gamma\right\rangle\right|^2 \\
&= \left|\sum_j U^*_{\beta j}U_{\alpha j}e^{-i\frac{m_j^2 L}{2E}}\right|^2 \\
&= \sum_j\sum_k U^*_{\beta j}U_{\alpha j}U^*_{\alpha k}U_{\beta k}e^{-i\frac{\Delta m^2_{jk}L}{2E}} \\
&= \sum_j |U_{\alpha j}|^2|U_{\beta j}|^2 + 2Re\left(\sum_j\sum_{k>j} U^*_{\beta j}U_{\alpha j}U^*_{\alpha k}U_{\beta k}e^{-i\frac{\Delta m^2_{jk}L}{2E}}\right) \\
&= \delta_{\alpha\beta} + 2Re\left(\sum_j\sum_{k>j} U^*_{\beta j}U_{\alpha j}U^*_{\alpha k}U_{\beta k}\left(e^{-i\frac{\Delta m^2_{jk}L}{2E}}-1\right)\right) \\
&= \delta_{\alpha\beta} + 2\sum_j\sum_{k>j} Re\left(U^*_{\beta j}U_{\alpha j}U^*_{\alpha k}U_{\beta k}\right)Re\left(\cos\frac{\Delta m^2_{jk}L}{2E} + i\sin\frac{\Delta m^2_{jk}L}{2E} - 1\right) \\
&\quad - 2\sum_j\sum_{k>j} Im\left(U^*_{\beta j}U_{\alpha j}U^*_{\alpha k}U_{\beta k}\right)Im\left(\cos\frac{\Delta m^2_{jk}L}{2E} + i\sin\frac{\Delta m^2_{jk}L}{2E} - 1\right) \\
&= \delta_{\alpha\beta} - 4\sum_j\sum_{k>j} Re\left(U^*_{\beta j}U_{\alpha j}U^*_{\alpha k}U_{\beta k}\right)\sin^2\frac{\Delta m^2_{jk}L}{4E} \\
&\quad + 2\sum_j\sum_{k>j} Im\left(U^*_{\beta j}U_{\alpha j}U^*_{\alpha k}U_{\beta k}\right)\sin\frac{\Delta m^2_{jk}L}{2E}
\end{aligned} \tag{1.22}$$

Note that $\Delta m^2_{jk} = m^2_j - m^2_k$ represents the difference in squared masses between eigenstates that drives neutrino oscillations. Trivially, if all eigenstates had the same mass then the transition probability would be $\delta_{\alpha\beta}$, allowing for no neutrino oscillations. In the derivation above, the unitarity of U was used in multiple places. It was used to split the double sum over j and k into a sum over the diagonal elements of U and a doubled sum of the upper diagonal elements of U, where only the real component survives since $U_{ab} = U^*_{ba}$. Additionally, it was used to introduce $\delta_{\alpha\beta}$:

$$\sum_j \sum_k U^*_{\beta j} U_{\alpha j} U^*_{\alpha k} U_{\beta k} = \left(\sum_j U^*_{\beta j} U_{\alpha j} \right) \left(\sum_k U^*_{\alpha k} U_{\beta k} \right) = \delta_{\beta\alpha} \delta_{\alpha\beta} = \delta_{\alpha\beta} \tag{1.23}$$

The oscillation probability for anti-neutrinos, $P_{\bar{\nu}_\alpha \to \bar{\nu}_\beta}$, can be computed by replacing U with U^* in Eq. 1.22. This also allows for CP violation to be computed as the difference between the neutrino and anti-neutrino oscillation probabilities:

$$P_{\nu_\alpha \to \nu_\beta} - P_{\bar{\nu}_\alpha \to \bar{\nu}_\beta} = 4 \sum_j \sum_{k>j} Im \left(U^*_{\beta j} U_{\alpha j} U^*_{\alpha k} U_{\beta k} \right) \sin \frac{\Delta m^2_{jk} L}{2E} \tag{1.24}$$

Referring back to the decomposition of U in Eq. 1.19 into its four degrees of freedom, it can be seen that δ_{CP} is the only variable capable of producing imaginary terms in U, and thus controls the degree of CP violation.

Before moving on, it is worth emphasizing some of the key features of neutrino oscillations introduced in the preceding derivations. Under the three-flavor standard model description, neutrino mixing is described through the PMNS matrix U by four parameters: three mixing angles and a complex phase that describes the degree of charge-parity violation. Neutrino oscillations between flavors also depends on the difference in mass (squared) between eigenstates, adding three additional parameters of interest (although it will later be discussed how two masses are so similar as to reduce the description to two measurable mass differences). The situation is further complicated by the dominant oscillation term's

dependence on the square of the mass difference, leaving it insensitive to the mass ordering. The remaining sensitivity directly available in the oscillation probability only appears if U has an imaginary component and correspondingly there is charge-parity violation in the neutrino sector. Alternatively, matter effects can be used to generate a potential that shifts the ν_e mass eigenstate value, revealing the hierarchy depending on whether the mass difference shrinks or grows, especially when compared against $\overline{\nu}_e$ measurements that would observe an opposite shift [21].

For all these parameters of interest, there are only two variables that can be directly controlled in experimental design. Neutrino energy can be controlled in accelerator experiments through the beamline used, and otherwise varies significantly between sources of cosmic neutrinos, solar neutrinos, atmospheric neutrinos, and neutrinos created in nuclear reactors. The more easily controlled variable is the distance travelled between a neutrino source and a detector, making detector placement a key design consideration in an experiment. However, these two variables only appear together as the ratio L/E, reducing the number of controllable degrees of freedom to one. As a result, the PMNS matrix has only been determined as well as it has through clever experiment designs that aim to isolate elements of the mixing matrix or other parameters of interest, as future sections will describe in more detail. The neutrino sector is not yet fully measured, however, as there remain unknowns such as the mass hierarchy [21], the value of δ_{CP} [22], the octant of θ_{23} [23], and the overall neutrino masses [24]. Next generation experiments are currently being constructed to determine answers for each of these unknowns [25].

1.4 Two-Flavor Oscillations

To help make sense of the oscillation probability formula, it can be useful to consider the simplified case of two-flavor oscillations. This is equivalent to assuming that two of the

mixing angles, say θ_{13} and θ_{23}, are 0, simplifying the PMNS matrix to:

$$U_{2\times 2} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \tag{1.25}$$

where ν_3 has been dropped as it does not mix, and the single remaining mixing angle is simply given as θ. The oscillation probabilities can then be computed to give the straightforward expressions:

$$P_{\nu_\alpha \to \nu_\alpha}(L) = 1 - \sin^2 2\theta \sin^2 \frac{\Delta m^2 L}{4E} \tag{1.26}$$

$$P_{\nu_\alpha \to \nu_\beta}(L) = \sin^2 2\theta \sin^2 \frac{\Delta m^2 L}{4E} \tag{1.27}$$

where α and β are presumed to be different flavors and the single remaining mass difference is given as Δm^2. This formula shows that there is a maximum degree of mixing established by the mixing angle, which is periodically achieved as the neutrino oscillates over its length scale $L_0 = \frac{4E}{\Delta m^2}$. This simplified two-flavor case may seem too basic to be of any real use in a three-flavor world; however, the particular values of mixing and mass splitting parameters allow for multiple situations to reduce to the two-neutrino case under reasonable approximations.

2 Neutrino Interactions

2.1 Importance of Cross Section Measurements

Accurate knowledge of neutrino cross sections is important for any neutrino experiment, especially modern accelerator oscillation experiments. These aim to determine the remaining unspecified neutrino parameters and explore new physics through precision neutrino appearance and disappearance searches. To achieve these goals, measurements must be highly accurate with low overall uncertainties. A disappearance measurement can be expressed in terms of the oscillation probability $P_{\nu_\alpha \to \nu_\alpha}$, the neutrino flux Φ, the total cross section σ, the selection efficiency ϵ, the detector response from true to reconstructed neutrino energy D, and the background event rate B:

$$N(E_\nu) = B + \int P_{\nu_\alpha \to \nu_\alpha}(E_\nu) \times \Phi(E_\nu) \times \sigma(E_\nu) \times \epsilon(E_\nu) \times D(E_\nu, E_{reco})\, dE_\nu \tag{2.1}$$

Clearly there are multiple factors that threaten to undermine the sensitivity of a measurement; experimentalists work hard to reduce sources of background while maintaining high (and well understood) selection efficiency, and uncertainties on the modeling of the flux and detector response are addressed through the use of a near detector.

By comparing the measurement at a near detector before significant neutrino oscillations to that at a far detector including oscillation effects, many uncertainties can be significantly reduced, including those on the flux modeling and detector response. This is because near and far detectors are often designed to be as similar as possible, so that a ratio of event rate measurements is insensitive to effects in common between the two detectors. As a result, cross section modeling can become the largest remaining source of uncertainty. To help address this situation, cross section measurements are produced at numerous experiments to aid the development of neutrino interaction models towards more accurate predictions with smaller uncertainties. This is the primary motivation of the triple-differential cross section

measurement in this dissertation.

Among cross section measurements, neutrino-energy-dependent cross section measurements are particularly useful in aiding oscillation searches. Neutrino energy is the most physically relevant quantity in a neutrino interaction; it is responsible for determining the overall interaction type. Identification of final state particles can help determine the interaction channel; however, Final State Interactions (FSIs) smear this mapping, reducing its usefulness. Neutrino energy provides an additional degree of separation, from Quasi-Elastic (QE) interactions at low energy to Deep Inelastic Scattering (DIS) at high energy. Furthermore, neutrino oscillations inherently depend on neutrino energy, so for oscillation experiments that observe a neutrino flux over a wide energy range it is important to both accurately reconstruct the neutrino energy and to understand the cross section and therefore interaction probability at that specific energy.

2.2 Neutrino Interaction Channels

Neutrinos have no electromagnetic or strong charge, so they interact with matter exclusively through the weak force carriers. These are the charged W and neutral Z bosons, giving rise to the corresponding Charged Current 2.2 (CC) and Neutral Current 2.3 (NC) interactions, respectively:

$$\begin{aligned} \nu_l A &\rightarrow l^- A^+ \\ \overline{\nu}_l A &\rightarrow l^+ A^- \end{aligned} \tag{2.2}$$

$$\begin{aligned} \nu_l A &\rightarrow \nu_l A \\ \overline{\nu}_l A &\rightarrow \overline{\nu}_l A \end{aligned} \tag{2.3}$$

Here $l \in \{e, \mu, \tau\}$ is used to represent the lepton flavor and A represents the nuclear target, such as an atom, nucleon, or quark. CC interactions produce final states that allow the neutrino flavor l to be easily identified, making them of particular interest to oscillation experiments.

Neutrino interactions can be further classified by the size of the nuclear target and the degree of interaction produced. Below 50 MeV neutrinos coherently scatter off entire nuclei in an elastic collision, producing a faint interaction signature of nuclear recoil that was only recently detected in 2017 by COHERENT [26]. As neutrino energy increases, the interaction becomes quasi-elastic, where a single nucleon (and the lepton) are ejected. While this definition seems simple, in practice it can be difficult to correctly identify QE interactions in cases where some final state particles are not able to be detected, such as the production and subsequent capture of a pion inside the nucleus [17, 27]. This is especially an issue with the larger nuclear targets used in many modern neutrino experiments, where there are more complicated nuclear effects to consider. As a result, FSIs can play an important role in interaction channel modeling and identification, and have been an increasingly large focus of cross section measurements. In addition to neutrino cross section measurements, high-statistics electron-nucleus scattering data has been a valuable resource in studying FSIs occurring within the atomic nucleus, despite the differences between neutrino and electron primary interactions [28].

QE interactions remain dominant up to $\sim 1\,GeV$, at which point there begins to be enough energy to excite the nucleon to resonant baryonic states. Most notably these include the Δ resonance with a mass of 1.232 GeV [29] and the Roper Resonance N^* with a mass of 1.44 GeV [30]. In both cases, the baryon is unstable and quickly decays, most commonly emitting a charged or neutral pion, but sometimes emitting multiple pions, a kaon or other meson, or a photon [31]. At these energies it is also possible for coherent inelastic neutrino-nucleus scattering, producing a pion with a particularly forward-scattered distribution [32]:

$$\nu_l A \rightarrow l^- A \pi^+$$

$$\nu_l A \rightarrow \nu_l A \pi^0$$
$$\overline{\nu}_l A \rightarrow l^+ A \pi^-$$
$$\overline{\nu}_l A \rightarrow \overline{\nu}_l A \pi^0 \tag{2.4}$$

Together, these interactions create a ~ 1 GeV landscape dominated by pion (and lepton) final states, with numerous other possible final states. This is all before considering final state interactions, which again serve to complicate an already varied landscape. Neutral pion production can be particularly relevant for measurements of other channels, as the π^0 produced does not interact electromagnetically and often cannot be detected until it decays, typically into a pair of photons [33]. If this decay is not spotted and associated with the neutrino interaction, the event becomes a background for QE interactions, and alternatively if the pion decay is detected but not properly identified, it can appear as an electromagnetic shower typical of an electron, such as in a ν_e appearance search. While baryonic resonance and coherent pion production are usually not the dominant interaction channels observed by an experiment, they still play an important role in the few GeV range.

Above ~ 3 GeV a significant fraction of interactions involve deep inelastic scattering, and above ~ 5 GeV DIS interactions dominate. Deep inelastic scattering involves high momentum transfer, often defined by requiring the four-momentum $Q^2 \equiv -q^2 > 1$ GeV, as well as high energy, requiring an invariant mass $W > 2$ GeV. In DIS interactions the neutrino can resolve individual quarks directly, with a resolution determined by the wavelength $\lambda = \hbar/\sqrt{Q^2}$. The high interaction energy allows for the nucleon to be broken apart, and as a result, final states can involve a large number of hadronic particles. DIS interactions have been studied at a number of experiments [34, 35, 36, 37, 38, 39], using various nuclear targets to produce measurements on the cross section as a function of neutrino energy, as well as on the weak mixing angle and structure functions that describe high energy neutrino-nucleus interactions. To first order, the DIS cross section increases linearly with neutrino energy, as predicted by

the quark parton model.

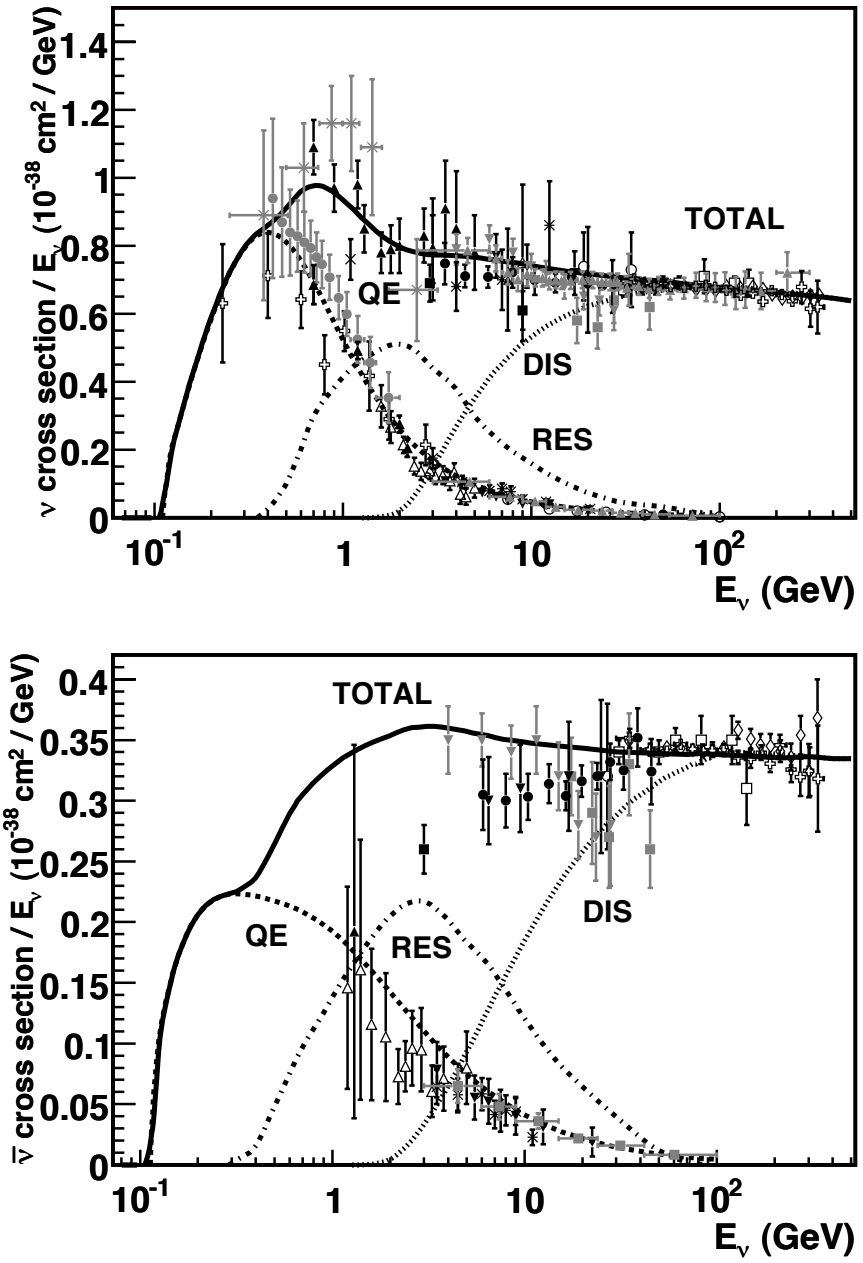

Figure 1: Total muon and antineutrino per nucleon CC cross sections divided by neutrino energy and plotted as a function of neutrino energy, taken from [31].

Clearly, the various neutrino interaction channels present a complicated picture for experimentalists to grapple with. It can be difficult to accurately identify final state particles and separate channels, especially when there can be overlapping visible interaction signatures

and FSIs are considered. The inclusive channel, consisting of all the different interaction channels as shown in Fig. 1, can help avoid these selection issues by making a selection agnostic to the underlying interaction channel. It also presents the largest possible statistics for a measurement, as no channels are excluded. For these reasons, it is typically the channel of choice for performing an oscillation measurement, and so is of particular interest for producing precision cross section measurements. However, there are challenges with reconstructing the inclusive channel: a wider signal definition means more phase space for backgrounds as well, the analysis needs to be able to reconstruct a wide range of topologies, and the modeling of exclusive channels is still important to prevent large systematic errors through mis-modeling.

2.3 Nuclear Effects

A crucial part of neutrino cross section modeling is bridging the gap between neutrino interactions with free nucleons imagined in a vacuum, and the complex reality of the many-body atomic nucleus. While the free-nucleon approximation may function well for simple atomic targets like Hydrogen, it neglects many nuclear effects in the heavy atomic targets more frequently seen in modern experiments. To understand where and how cross section measurements can aid model development, it can help to see what problems these models have to address. These can be split into two groups: the modeling of the initial state of the nucleus and its nucleons, including their kinematic distributions and bound states, and the modeling of the passage of final state particles through the nucleus following an interaction.

Modeling neutrino-nucleus interactions begins with a description of the nucleons' initial states, achieved succinctly through the use of a spectral function $P(k,E)$, which gives the probability that a nucleon struck in a neutrino-nucleus interaction will have initial momentum k and removal energy E [40]. The Relativistic Fermi Gas (RFG) model produces a spectral function through a simple description of nucleons as a degenerate gas of non-interacting

particles up to the Fermi momentum k_F [41]:

$$P(k,E) = \delta\left(E + \frac{\hbar^2}{2m}k^2\right) \tag{2.5}$$

This description is expanded on in the Local Fermi Gas (LFG) model, which treats the Fermi momentum as a function of radius through the nucleon density function $\rho(r)$ [42]:

$$k_F(r) = \left(\frac{3}{2}\pi^2\rho(r)\right)^{1/3} \tag{2.6}$$

Perhaps the simplest model of how a neutrino interacts with the nucleus is the Impulse Approximation (IA), which assumes that the neutrino interacts with individual nucleons and that final states produced in this interaction do not interact with the remaining nucleons in the atom [43]. This ignores the nucleon-nucleon correlations in the initial state (as well as FSIs after interaction), so as a result models attempt to describe these effects separately and add them back in. Nucleon-nucleon correlations serve to form bound states within the nucleus of an atom, through quarks and gluons at short distances, and through pion(s) at longer distances [44], earning them the name Meson Exchange Current (MEC). These bound states can affect the distribution of final state topologies, such as through the correlated ejection of protons and neutrons [45]. One noteworthy effect is the two-nucleon correlation, called the Two-Particle-Two-Hole (2p2h) interaction [46], as it is the lowest order bound state. Long distance interactions are particularly important at low and medium energies ($Q^2 <$ 1 GeV), and are described using the Random Phase Approximation (RPA), which condenses the large number of small interactions between distant nucleons into an overall effective potential [47]. The nuclear initial state can be probed experimentally through the study of Transverse Kinematic Imbalance (TKI), which is driven by the kinematics of the nuclear initial state [48, 49]. In total, there is significant focus on accurate modeling of the nuclear initial state, which makes sense given the large impact it can have on final states and kinematic

distributions produced.

In addition to nuclear initial state modeling, significant attention is given to an accurate treatment of final state particles as they travel through the nucleus. A simple but useful approach is through the hA model, which skips the many interactions between a hadron and the nucleons and instead uses the total cross section for each nuclear process [50]. Although this reductive approach has a limited capacity to describe the nuances in nature, it achieves decent accuracy from a wealth of hadron-nucleus cross section data while allowing for fast computation and simple re-weighting.

Going beyond the simple hA model, the Intranuclear Cascade (INC) model provides a description that computes a large number of hadron-nucleon interactions [51, 52]. This approach simulates interactions probabilistically over steps of length L based on the mean free path $\lambda = \frac{1}{\rho\sigma}$, computed from the nucleon density ρ and the hadron-nucleon cross section σ [53]:

$$P(L) = e^{-L/\lambda} \tag{2.7}$$

One such INC model is the hN implementation, used in event generators such as Genie v2.6. Owing to its more detailed modeling, the hN implementation is able to more accurately describe final state interactions than hA in Genie [50]. The modeling of FSIs is difficult to do accurately, and event generator improvements are continuously made [54, 55], aided by cross section measurements.

3 History of Neutrino Experiments

3.1 Discovery

There have been many neutrino experiments performed since the particle's conception by Wolfgang Pauli in 1930. They can largely be grouped into four sources of neutrino emissions: solar, atmospheric, reactor, and accelerator, as well as four main methods of detection: nuclear emulsion, Cherenkov radiation, scintillation light, and Liquid Argon Time Projection Chambers (LArTPC).

Solar and atmospheric neutrino experiments have led the field to new physics through the latter half of the 20th century, first with hints of neutrino oscillations from the Homestake experiment [56], and then conclusive evidence of their existence with Super-Kamiokande [57, 58] and Sudbury Neutrino Observatory (SNO) [59]. Solar electron neutrinos are produced in the sun as a byproduct of pp-chain fusion reactions, predominantly through the formation of deuterons and sub-MeV neutrinos, with branches involving ^{7}Be and ^{8}B contributing the majority of higher energy neutrinos [60]:

$$p + p \rightarrow {}^{2}_{1}\mathrm{D} + e^{+} + \nu_e \tag{3.1}$$

$${}^{7}_{4}\mathrm{Be} + e^{-} \rightarrow {}^{7}_{3}\mathrm{Li} + \nu_e \tag{3.2}$$

$${}^{8}_{5}\mathrm{B} \rightarrow {}^{8}_{4}\mathrm{Be} + e^{+} + \nu_e \tag{3.3}$$

The presence of electrons in matter alters the Hamiltonian describing neutrino flavor time evolution. Weak interactions between electrons and sufficiently energetic (>1MeV) electron neutrinos, known as the MSW effect, or matter effect [61, 62], create a potential experienced by electron neutrinos. The mass eigenstate and thus the oscillation probabilities are changed

as a result, as these depend on the mass difference (squared) between eigenstates, as shown in Eq. 1.22. As a result, although the sun produces overwhelmingly electron neutrinos, the flux leaving the sun at high energies such as through the ^{8}B channel is significantly mixed between the three flavors, with a P_{ee} survival probability of 34%. Low energy neutrinos such as from the deuteron branch are not effected by the MSW effect, and reach the earth with a P_{ee} survival probability of 60% due to vacuum oscillations between the sun and earth.

The first hint of neutrino oscillations was detected by the Homestake experiment in the late 1960s. Located deep underground in the Homestake gold mine in Lead, South Dakota, the experiment counted the rate of neutrino capture on ^{37}Ar through the measurement of the resulting ^{37}Ar:

$$\nu_e + {}^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Ar} + e^- \tag{3.4}$$

It wasn't until 30 years later with the measurements from Super-Kamiokande and SNO that neutrino oscillations were conclusively proven. Kamiokande II and later Super-Kamiokande are water Cherenkov detectors [63], which observe the Cherenkov radiation generated by fast-moving electrically charged particles that exceed the speed of light in water. Kamiokande II was able to reconstruct the direction of incident neutrinos from the Cherenkov light cone generated by high-energy electrons produced by electron neutrino charged current interactions [64]:

$$\nu_e + X \rightarrow X' + e^- \tag{3.5}$$

This confirmed the sun as the source of the neutrinos observed in Homestake and Kamiokande II.

Furthermore, Kamiokande II also measured a rate deficit of neutrino interactions when compared to predicted values using the Standard Solar Model (SSM) [65], affirming the deficit seen in Homestake. Super-Kamiokande went further, reconstructing atmospheric muon neutrinos generated from the decay of pions created by cosmic rays. By reconstructing the direction of the incoming muon neutrinos, Super-Kamiokande demonstrated the existence of

neutrino oscillations by measuring a relative rate deficit in events entering the detector from below, caused by oscillations into tau neutrinos over the longer distance travelled through the earth. Furthermore, adding the tau-like-event measurement to the muon measurement accounted for the muon deficit, demonstrating the muon-to-tau oscillation directly.

Finally, in 2001 SNO provided definitive evidence of electron neutrino oscillations causing the solar electron neutrino deficit. Similar to Super-Kamiokande, SNO is a Cherenkov detector that measured the electron neutrino flux from the electrons produced in the CC interaction 3.1, observing the familiar deficit. Additionally, by measuring the neutral current interaction rate, which was independent of neutrino flavor, the total neutrino flux was measured and found to agree with prediction:

$$\nu_l + D \rightarrow p + n + \nu_l \tag{3.6}$$

Therefore, the deficit of electron neutrinos but agreement in overall neutrino flux showed that electron neutrinos must be oscillating to other flavors.

3.2 Oscillation Parameter Searches

While solar and atmospheric neutrino experiments are responsible for discovering the existence of neutrino oscillations, they are not well equipped to measure all of the oscillation parameters. Probing a given parameter requires an experimental setup within a specific L/E range, the appropriate source neutrino flavor, and meeting energy constraints to enable the detection of the particular neutrino interaction of interest. These requirements come together to motivate human-produced neutrino sources, where the length scale can be freely set, and, in the case of accelerator experiments, the energy scale can even be chosen.

KamLAND is a reactor neutrino experiment that demonstrates the power in being able to choose the neutrino oscillation baseline distance. The liquid scintillator [66] detector is located in Japan near Kamiokande, allowing it to see $\bar{\nu}_e$ from a total of 55 Japanese nuclear

reactor cores, with a weighted average baseline of 180 km. Given the solar Δm^2_{12} estimate of $7.5 \times 10^{-5} eV^2$, the $\nu_e \rightarrow \nu_\mu$ oscillation wavelength/E_ν is 33 km/MeV. KamLAND estimated the incident neutrino energy from the positron produced, and observed two full oscillation periods over the measured energy range, with best fits improving estimates of Δm^2_{12} and θ_{12}, in particular resolving the degeneracy of θ_{12} around 45° and confirming the Large Mixing Angle (LMA) scenario [67].

Going beyond θ_{12} and θ_{23}, a precision measurement of the mixing angle θ_{13} used in $\nu_e \rightarrow \nu_\tau$ is beyond the scope of solar and atmospheric experiments. The best way to measure θ_{13} without involving θ_{12} or θ_{23} (and their uncertainties) in the calculation is through $\nu_e \rightarrow \nu_e$ or $\overline{\nu}_e \rightarrow \overline{\nu}_e$ disappearance. At an L/E of roughly 0.5 km/MeV, the disappearance probability from $\overline{\nu}_e \rightarrow \overline{\nu}_\mu$ is negligible. However, the atmospheric Δm^2_{23} estimate of $2.4 \times 10^{-3} eV^2$ predicts maximal $\overline{\nu}_e \rightarrow \overline{\nu}_\tau$ oscillations, measuring θ_{13} via:

$$\begin{aligned} P_{\overline{\nu}_e \rightarrow \overline{\nu}_e} &= 1 - \cos^4\theta_{13} \sin^2 2\theta_{12} \sin^2 \frac{\Delta m^2_{12} L}{4E} \\ &\quad - \sin^2 2\theta_{13} \left(\cos^2\theta_{12} \sin^2 \frac{\Delta m^2_{13} L}{4E} + \sin^2\theta_{12} \sin^2 \frac{\Delta m^2_{23} L}{4E} \right) \\ &\approx 1 - \sin^2 2\theta_{13} \sin^2 \frac{\Delta m^2_{23} L}{4E} \end{aligned} \tag{3.7}$$

where the approximation $\sin^2 \frac{\Delta m^2_{12} L}{4E} \approx 0$ was used to drop the first term, and $\Delta m^2_{13} \approx \Delta m^2_{23}$ was used to simplify the second to fully eliminate θ_{12} dependence through $\sin^2\theta_{12} + \cos^2\theta_{12} = 1$.

Reactor neutrino experiments are ideal for this search; the $\overline{\nu}_e$ produced allow for $\bar{\nu}_e \rightarrow \bar{\nu}_e$ disappearance to be measured. The Daya Bay [68], Double CHOOZ [69], and RENO [70] experiments all performed this measurement with scintillator detectors in the early 2010s, generating similar results for a global fit of $\theta_{13} = 8.93^\circ \, ^{+0.34}_{-0.45}$ [71, 72]. Daya Bay led the series of measurements, featuring three different detector and reactor locations for weighted baseline lengths of 360m, 500m, and 1650m. By comparing the measurements of the near detectors with those of the far detector, a clear $\bar{\nu}_e$ deficit (before considering oscillations) at the far

detector was observed, leading to the θ_{13} best fit measurement.

Thus far, solar, atmospheric, and reactor neutrino experiments have been explored, and shown to allow for a wide range of available baselines to effectively measure the different mixing and mass difference parameters. However, only the atmospheric neutrinos exceed a few MeV in energy, and these still do not probe the highest energies. Exploring neutrino interactions at high energies can be particularly important for probing the large-mass-splitting regime for oscillations, as well as for direct detection of ν_τ through the production of τ particles, which require GeV energy scales due to its mass. Looking ahead, access to high energy neutrinos can also be useful for experiments such as DONUT [73] and DUNE [74] that aim to reconstruct ν_τ interactions.

The motivation to choose the neutrino energy regime in an experiment is clear, and particle accelerators provide the means to do so. In a particle accelerator, protons are accelerated using electromagnetic fields to highly relativistic speeds and smashed into a target. Typically these collisions create charged π and K mesons that are focused into a tight beam via a magnetic horn before creating neutrinos as they decay at rest. A significant amount of the incident proton's momentum is transferred to the meson and then the neutrino, causing it to decay within a cone of a few degrees in width along the beam direction set by the incoming proton beam and meson focusing horn. Any other decay products are stopped by matter along the beamline before reaching a neutrino detector, leaving a pure neutrino beam incident on the neutrino detector.

The T2K experiment [75] uses the freedom allowed in experimental design of accelerator experiments to simultaneously measure multiple neutrino parameters. A 99.5% pure ν_μ beam is produced in Tokai, Japan, and sent 295 km to the Super-Kamiokande detector discussed previously. The neutrino beam reaches multiple GeV in energy, but is purposefully directed so that the detector is 2.5° off-axis, changing the flux profile to sharply peak at 600 MeV. This achieves the familiar L/E of 0.5 km/MeV to allow for maximal $\nu_\mu \rightarrow \nu_\tau$ oscillation,

Table 1: Best fit values with uncertainties for each neutrino mixing and mass splitting parameter. Source [78].

Parameter	Best Fit Value $\pm 1\sigma$
θ_{12} (degrees)	$33.48^{+0.77}_{-0.74}$
θ_{13} (degrees)	$8.52^{+0.20}_{-0.21}$
θ_{23} (degrees)	$42.2^{+0.1}_{-0.1}$ or $49.4^{+1.6}_{-2.0}$
δ_{CP} (degrees)	251^{+67}_{-59}
$\Delta m^2_{21} \times 10^{-5}$ (eV)	$7.50^{+0.19}_{-0.17}$
(normal) $\Delta m^2_{32} \times 10^{-3}$ (eV)	$+2.458^{+0.046}_{-0.047}$
(inverted) $\Delta m^2_{32} \times 10^{-3}$ (eV)	$-2.448^{+0.047}_{-0.047}$

allowing T2K to produce high precision measurements of θ_{23} and Δm^2_{23} [76]. Additionally, the T2K beamline is so pure in ν_μ that a ν_e appearance search is possible. This allows for a measurement of θ_{13} [77] via:

$$P_{\nu_\mu \to \nu_e} \approx \frac{1}{2} \sin^2 \theta_{13} \sin^2 \frac{\Delta m^2_{23} L}{4E} \tag{3.8}$$

Counter-intuitively, the $\nu_\mu \to \nu_e$ oscillation probability does not (to good approximation) depend on θ_{12} at short distances. This is because Δm^2_{12} is so small and thus the oscillations require a much larger distance; meanwhile Δm^2_{23} and Δm^2_{13} are similarly large, so that the oscillation can be thought of as primarily $\nu_\mu \to \nu_\tau \to \nu_e$ rather than the more straightforward $\nu_\mu \to \nu_e$ oscillation. T2K was able to measure 28 ν_e with an expected background of only 11 events, constituting a positive measurement of $\nu_\mu \to \nu_e$ oscillations at 7σ and contributing to the global θ_{13} fit described earlier. Table 1 shows the best fit values for each of the neutrino mixing and mass splitting parameters.

Particle accelerator experiments have also been able to confirm the $\nu_\mu \to \nu_\tau$ oscillations first found in atmospheric experiments. The DONUT experiment [73] first detected the ν_τ using the Fermilab Tevatron to generate the high energy beam needed, and used a nuclear emulsion detector to identify τ particles indicative of ν_τ CC interactions through the tau's

kinked track at the point of decay. Building on this, the OPERA experiment [79] at CERN used a similar emulsion detector to identify ν_τ interactions originating from a ν_μ beam 732 km away in Gran Sasso, providing direct evidence of $\nu_\mu \rightarrow \nu_\tau$ oscillations.

Because of their high energies and freedom to choose very short baselines, accelerator experiments are able to pick L/E values that probe drastically different mass splittings than the measured $\Delta m^2_{12} \approx 7.5 \times 10^{-5} eV^2$ and $\Delta m^2_{23} \approx 2.4 \times 10^{-3} eV^2$. LSDN was a scintillator detector built in the 1990s before the three mass splittings were well measured, and used an extremely short baseline of 30m to probe the $\Delta m^2 \approx 1 eV^2$ region. Surprisingly, it measured a ν_e appearance, indicating $\nu_\mu \rightarrow \nu_e$ oscillations [80], a result which is in conflict with the numerous experiments discussed above, as well as cosmological measurements that set a hard limit of three weakly interacting neutrinos [81]. The three neutrinos discussed thus far, ν_e, ν_μ, and ν_τ, allow only two unique mass splittings between them, so a third unique mass splitting at $1eV^2$ would require a fourth neutrino. Cosmological constraints disallow this new particle from interacting weakly, so the hypothetical particle is called the sterile neutrino. As a sterile neutrino is not predicted by the standard model, it falls into the category of proposed physics Beyond the Standard Model (BSM). Of course there are other potential explanations for the ν_e excess observed by LSND, BSM and otherwise; further measurements are needed.

The MiniBooNE experiment was built to investigate the LSND anomalous measurement using the Fermilab Booster Neutrino Beam (BNB) and a Cherenkov detector. Using a different beamline, energy range, and detector type, MiniBooNE would represent a fully independent measurement to corroborate or reject the LSND data. MiniBooNE did find an anomalous excess in ν_e appearance, particularly at low energy below 600 MeV, now dubbed the Low Energy Excess (LEE). Together, the MiniBooNE and LSND measurements represent a 6.1σ disagreement with the standard model [82]. Either there are significant errors in the modeling used in these experiments, such as in the rate of background pion decays appearing as electrons, or there is some new physics explanation such as a sterile neutrino. A sterile

neutrino could explain the excess through the oscillation $\nu_\mu \rightarrow \nu_s \rightarrow \nu_e$ adding to the small rate of direct $\nu_\mu \rightarrow \nu_e$ oscillations predicted.

3.3 Overview of Detector Types

Investigating the Low Energy Excess of MiniBooNE is one of the primary goals the MicroBooNE experiment [83]. To explain how MicroBooNE can provide new information despite using the same Fermilab BNB beamline as MiniBooNE and a very similar L/E, 470 m instead of 500 m, the differences in detector types must be discussed. As mentioned earlier, most experiments use either nuclear emulsion, Cherenkov radiation, scintillation light, or LArTPCs to detect neutrino interactions. Nuclear emulsion uses a photographic plate coated in photographic emulsion, capable of detecting high energy ionizing particles with sub-micrometer resolution [84]. By stacking plates, a full 3D image of particle trajectories can be created. This high resolution imaging was particularly helpful for the DONUT and OPERA experiments in identifying the complex event topologies of energetic tau neutrino interactions. However, the emulsion approach has a few drawbacks. First, although the positional resolution is spectacular, there is no direct calorimetric information available for use in determining the energies involved. Additionally, the technology struggles to scale well; although a single plate of emulsion may not be too expensive to purchase or later to analyze for data, the technology struggles to maintain practicality for use in modern kiloton+ size detectors.

Not mentioned thus far are bubble chambers, which face similar advantages and drawbacks to emulsion detectors. The detector volume is filled with superheated liquid which vaporizes in the presence of highly energetic particles, creating a precise record of the particle trajectory when photographed by cameras around the detector [85]. Similar to emulsion detectors, they face difficulties in scaleability, both in analyzing the images produced, and in maintaining the large volume of superheated liquid, often liquid hydrogen.

By comparison, Cherenkov and scintillation detectors are scaleable, making them the

chosen detector type for many of the experiments discussed. As discussed earlier, Cherenkov detectors utilize the reduced speed of light in a medium such as water to produce Cherenkov radiation. This light travels outward in a cone coaxial with the direction of motion of the radiating particle, and is detected by photo-multiplier tubes (PMTs). This allows the particle's trajectory and point of origin to be roughly determined, but with far less accuracy than emulsion detectors. As a result, it can be difficult to distinguish between different particle signatures, such as electromagnetic showers originating from electrons or photons. However, Cherenkov detectors can be built to be extremely large, such as in the case of Super-Kamiokande, and are able to reconstruct the energy of the radiating particle involved.

Scintillation detectors produce light through a different mechanism than Cherenkov detectors, but similarly detect it with PMTs. Scintillating material throughout the detector can become energized by the passage of ionizing particles. The scintillating material will then emit photons of its own after a short period of time, generating a trail of where the ionizing particle traveled. Like Cherenkov detectors, scintillation detectors are capable of measuring the energy of the ionizing particles involved, aiding in the physics analysis. It is even possible to combine these detector technologies, such as in LSND, which measures the Cherenkov and scintillation light produced. However, the spacial resolution of scintillation detectors is still far inferior to that of emulsion detectors.

There is a clear trade-off between the precision of emulsion and bubble chambers, and the scaleability and calorimetry of scintillators and Cherenkov detectors. As a relatively new detector type to neutrino physics, the LArTPC offers many of the benefits of each of these detectors all in one. High energy ionizing particles, such as those produced from neutrino interactions, leave a trail of detached electrons as they travel through the liquid argon detector [86]. An applied electric field across the detector draws the electrons towards the cathode, where multiple (often three) stacked wire planes detect the presence of passing electrons through current induced on the wires. The wires of each plane are arranged at

an angle to the other planes, so that while a single plane only provides information along one spatial direction, combined they can give an accurate 2D picture. The initial neutrino interaction time is known through the use of PMTs which observe the prompt light flash generated. The interaction time is combined with the time information of the induced wire-plane currents to project the charge measurements to a 3D location based on the known drift velocity.

The 3D images produced by LArTPCs are not quite as accurate as emulsion detectors, but still offer millimeter level resolution, which makes is possible to clearly identify many particle topologies and decays. To add to this, the number of electrons ionized is proportional to the energy deposited, meaning that the LArTPC is a highly effective calorimeter. This can be useful for big picture physics goals such as measuring the total neutrino energy, as well as for smaller tasks such as identifying particles and their directionality by measuring the rate of energy deposition along a particle track. LArTPCs are also practical for large-scale experiments, as the argon required is not overly expensive, and the number of wires in a plane (and associated electronics) scales with the detector length not the detector volume. As a relatively new detector type, perhaps their largest potential drawback is the challenge posed in fully reconstructing highly accurate 3D images from the raw current waveform recorded. However, this is a task that has now been accomplished to a high degree of success, paving the way for future analyses.

4 The MicroBooNE Experiment

4.1 Primary Goals

The MicroBooNE experiment was designed with three primary goals: investigate the LSND and MiniBooNE Low Energy Excess (LEE), provide research and development experience on building large scale LArTPCs, and produce a wide range of neutrino-argon cross sections. While the MiniBooNE LEE measurement achieved a high statistical significance, it struggled to produce a pure event selection, as $NC\,\pi^0$ events dominated the ν_e signal, shown in Fig. 2. Because of the high background, mis-modeling of the π^0 background prediction beyond its uncertainty could have a significant impact on the measured tension with the standard model. The MicroBooNE experiment took data over 5 years, or runs, from 2015 to 2021, and addresses this ambiguity by using a LArTPC detector. Unlike the MiniBooNE Cherenkov detector, it is able to differentiate electron showers from the $\pi^0 \rightarrow \gamma\gamma$ decay. MicroBooNE investigated the LEE region under both $CC\ \nu_e$ and $NC\ 1\gamma$ hypotheses, disfavoring both explanations as sole explanations for the MiniBooNE LEE [83, 87]. This leaves the interpretation of the MiniBooNE LEE unclear, with candidate explanations including potential mis-modeling as described earlier, potential BSM interpretations beyond the single sterile neutrino case, or a combination of multiple factors.

MicroBooNE has also contributed to the field of neutrino experiments through the insight gained in building and running the experiment. The field of accelerator neutrino physics has embraced the LArTPC detector in recent years, with all three Short Baseline Neutrino (SNB) program experiments and the upcoming DUNE experiments using LArTPC detectors. In this context, the MicroBooNE experiment is an important pioneer of the technology. MicroBooNE is the first large-scale LArTPC to use cold electronics to significantly reduce the noise introduced in wire-plane readouts. Additionally, MicroBooNE dealt with multiple unexpected challenges such as dead wires across the three wire planes. As a result of encountering these

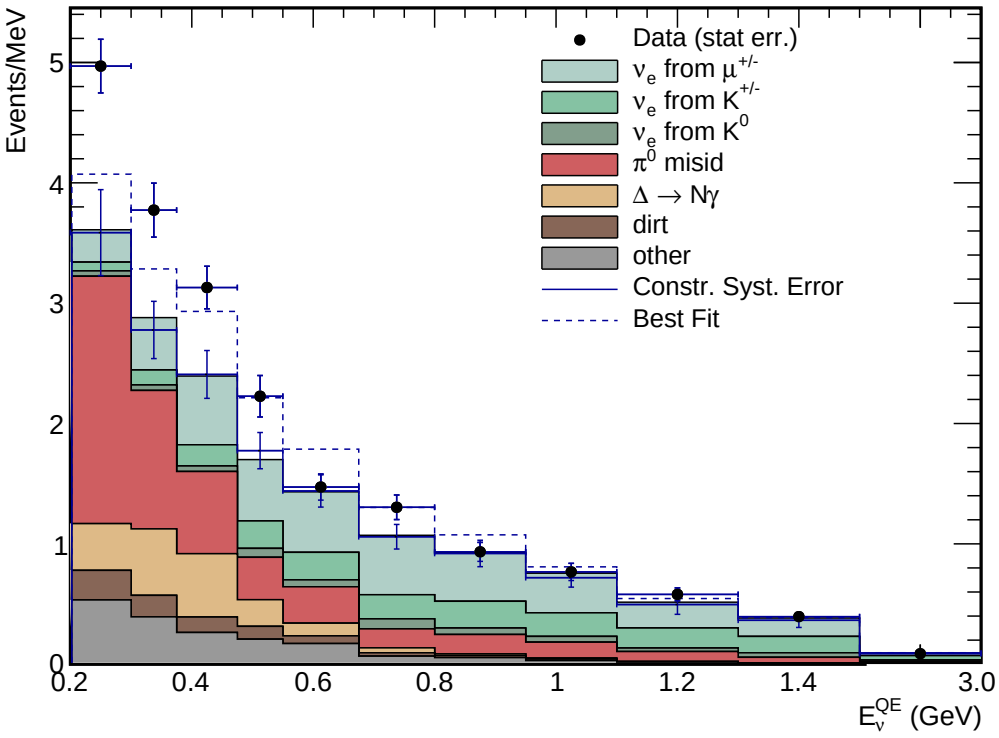

Figure 2: The MiniBooNE $CCQE$ ν_e measurement over E_ν^{QE} using 12.84×10^{20} POT, as published in [82]

issues, future LArTPC design has been improved to greatly reduce their risk of occurrence, or otherwise mitigate the challenges they present. On the software side, the development of MicroBooNE analyses has resulted in numerous advances to the field that can be applied to similar tasks in future LArTPC experiments.

To date, MicroBooNE has collected the largest data set of neutrino-argon interactions, enabling a large suite of cross-section measurements. Key to these measurements is the LArTPC detector technology, featuring precision spacial resolution and calorimetry to enable strong particle identification capabilities. In turn, this enables the reconstruction and high-efficiency event selection for many topologies. Furthermore, MicroBooNE can leverage its long-term exposure to both the BNB and NuMI beamlines at Fermilab, granting high statistics for both electron and muon neutrinos. As discussed earlier, these cross section measurements help advance the field by guiding the path for model development. Moreover,

since MicroBooNE analyses report cross sections on argon, the same detection medium used in SBND, ICARUS, and DUNE, the cross section measurements at MicroBooNE are especially valuable to the suite of upcoming US-based accelerator neutrino experiments.

4.2 The Booster Neutrino Beam

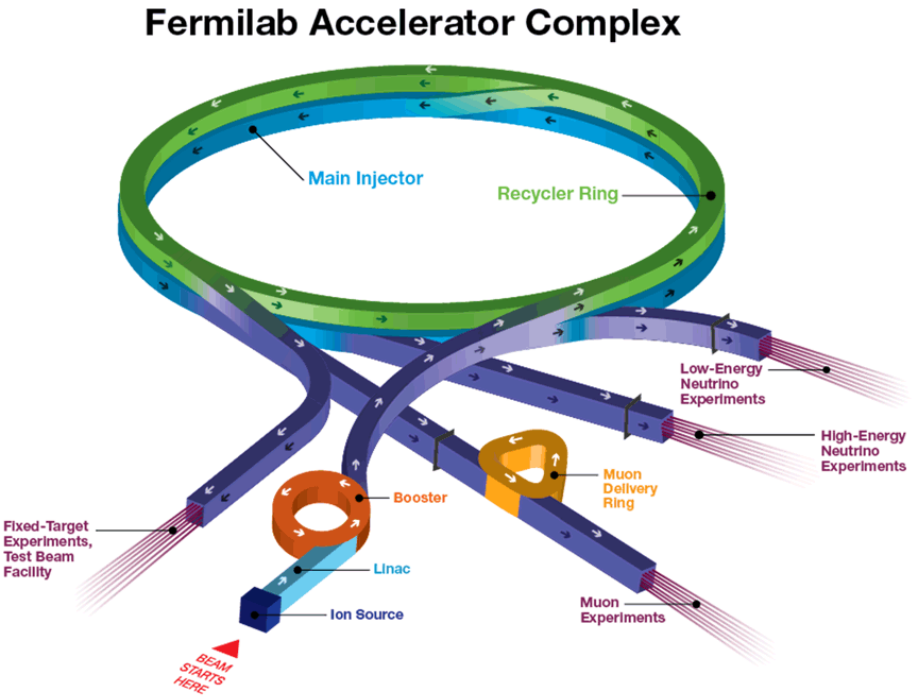

Figure 3: The Fermilab accelerator complex. MicroBooNE uses the Booster Neutrino Beam, labeled here for Low-Energy Neutrino Experiments.

The MicroBooNE experiment is situated along the Booster Neutrino Beamline [88] at Fermilab. The Fermilab accelerator complex consists of the Ion Source, Linear Accelerator (Linac), Booster ring, and in the case of the Neutrino Main Injector (NuMI) beamline, the Main Injector and Recycler, as shown in Fig. 3. The neutrino beam is produced through the acceleration of protons in several stages. First H^- ions are accelerated to 750 keV at the Ion

Source before reaching 400 MeV in the Linac. Before entering the Booster, the electrons are stripped off with a carbon foil, leaving a proton beam. The Booster is a 474 m circumference synchrotron capable of accelerating protons to 8 GeV. A kicker magnet creates 1.6 μs spills of $\sim 4 \times 10^{12}$ protons at a time to be sent to the Booster Neutrino Beamline. Two toroids measure the beam flux with a 2% margin of error, and magnets focus the beam into a $\sim$ 1 mm wide column through a beryllium target.

The MicroBooNE flux prediction uses the flux prediction for MiniBooNE [89] updated to the MicroBooNE detector location [90]. The p-Be interactions at the target produce mainly $\pi^{\pm}$ and $K^{\pm}$, but also include p, n, and K^0_L. The target is located inside a toroid called the horn, which uses its magnetic field to focus particles of the preferred sign, and de-focus particles of the opposite sign. When run in Forward Horn Current (FHC) mode, it focuses positively charged particles which later decay to neutrinos in a 50 m long decay pipe filled with air, while alternatively in Reverse Horn Current (RHC) mode it focuses negatively charged particles to produce anti-neutrinos on decay:

$$\begin{aligned} \pi^+ &\rightarrow \mu^+ \nu_\mu \\ \pi^- &\rightarrow \mu^- \overline{\nu}_\mu \end{aligned} \tag{4.1}$$

$$\begin{aligned} K^+ &\rightarrow \mu^+ \nu_\mu \text{ or } \pi^+ \pi_0 \\ K^- &\rightarrow \mu^- \overline{\nu}_\mu \text{ or } \pi^- \pi_0 \end{aligned} \tag{4.2}$$

Notably, muon decay produces different neutrinos than the type selected by the FHC or RHC mode, constituting a contamination to both oscillation measurements through intrinsic ν_e in the beamline, and cross section measurements through an (anti-)neutrino background:

$$\mu^+ \rightarrow e^+ \overline{\nu}_\mu \nu_e$$

$$\mu^- \rightarrow e^- \nu_\mu \overline{\nu}_e \tag{4.3}$$

At the end of the decay pipe, a concrete absorber stops non-neutrino particles so that a pure neutrino beam reaches the MicroBooNE detector 470 m past the beryllium target. The neutrino flux at MicroBooNE is estimated to consist of 93.6% ν_μ, 5.68% $\overline{\nu}_\mu$, 0.52% ν_e, and 0.05% $\overline{\nu}_e$, shown per POT and as a function of neutrino energy in Fig. 4.

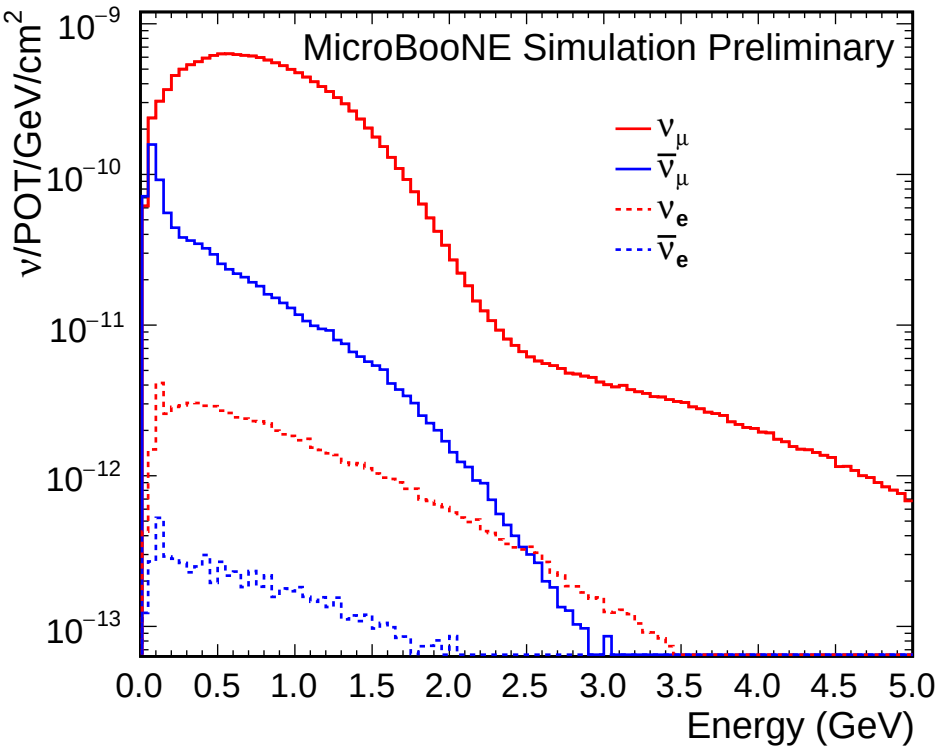

Figure 4: The absolute neutrino flux prediction through the MicroBooNE detector as calculated by the beam simulation. Shown is the flux for ν_μ, $\overline{\nu}_\mu$, ν_e, and $\overline{\nu}_e$ averaged through the TPC volume with dimensions 2.56 m (x) × 2.32 m (y) × 10.36 m (z). Source: Booster Neutrino Flux Prediction at MicroBooNE [82]

4.3 The MicroBooNE LArTPC

The MicroBooNE detector is a 2.56 m (x) × 2.32 m (y) LArTPC with an 85 tonne fiducial volume, encased in a cylindrical cryostat for a total mass of 170 tonnes of liquid argon. The cryostat

maintains a sealed environment of liquid argon (86−89 K) at 1.24 bar of pressure, and contains the cold electronics, field cage, cathode, and anode plane assemblies. A voltage of −70 kV is applied to the cathode plane to generate a nearly-constant electric field of 273.9 V/cm across the 2.56 m drift dimension (x axis). A series of three wire planes are located at the anode to detect the drift of electrons dislodged by ionizing particles, and an array of 32 PMTs are arranged behind the wire planes to detect prompt scintillation light. The detector is aligned along the neutrino beam (z axis) to present the longest possible dimension to image the trajectories of particles generated in neutrino interactions.

4.3.1 Liquid Argon Physics

The detector uses argon as the target mass for neutrino interactions for multiple reasons. First, as a noble gas, or rather noble liquid, it is chemically inert. This allows free electrons liberated by ionizing particles to drift through the argon medium without recombining with an argon atom. In fact, even slight impurities in the argon can drastically reduce the free electron lifetime. However, the MicroBooNE LAr purity achieved is quite high, allowing for a mean electron lifetime of 18 ms under the 273.9 V/cm electric field [91]. A long lifetime means that more electrons will reach the wire planes for a stronger signal and more reliable detection of energy deposited in the detector. Argon is used instead of other noble gases because it is both heavy, providing more target nucleons to generate a larger ν-Ar cross section, and relatively cheap to produce, owing in part to the fact that it is the third most abundant element in the atmosphere, comprising ~ 1%. Finally, argon acts as an excellent scintillator [92], which makes the detection and identification of neutrino interactions much easier through use of the light detection system.

In ν-Ar interactions a significant amount of energy is transferred to the final state particles. Energetic electrically charged particles ($\pi^{\pm}$, μ, p, e) ionize argon atoms as they travel through the argon medium, producing positive argon ions and free electrons. This ionization deposits

energy into the LAr, and the energy deposition per unit distance $\frac{dE}{dx}$ is well described by the Bethe-Bloch equation [17]:

$$\langle -\frac{dE}{dx} \rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \tag{4.4}$$

In this equation, K is a coefficient derived from the electron mass and volume (~ 0.307 MeV cm^2/mol), z represents the incident particle's charge (± 1), Z represents the atomic number (18 for Ar), A represents the atomic mass (~ 40 g/mol for Ar), γ represents the Lorentz factor, c represents the speed of light, β represents the velocity divided by c, m_e represents the electron mass, W_{max} represents the maximum possible energy transfer to an electron in a single collision, I represents the mean excitation energy, and $\delta(\beta\gamma)$ represents a density effect correction to ionization energy loss. Fig. 5 shows the energy deposition rate over the range of energies observed in MicroBooNE.

Of particular interest is the so-called “Minimum ionization" region, where the energy deposition rate is largely flat as a function of energy, and low compared to the overall particle energy ($< 1\%$). This means that particles in this region, called Minimumly Ionizing Particles (MIPs), can traverse a large distance while depositing a roughly constant amount of energy. Once an ionizing particle loses enough energy, its energy deposition rate increases, causing the particle to abruptly stop over a few centimeters as it deposits its remaining energy at a high rate, known as the Bragg peak. The minimumly ionizing region and Bragg peak are useful both for calibrating the detector response and identifying particle type and directionality. In contrast to the ease of ionizing particle detection, neutrally charged particles leave no trail through the detector, and can usually only be spotted upon decay (such as $\pi^0 \rightarrow \gamma\gamma$) if at all.

4.3.2 Recombination

Directly following the ionization of an electron, there is a relatively high chance of recapture by a nearby Ar^+ ion, called recombination. By representing the fraction of electrons that

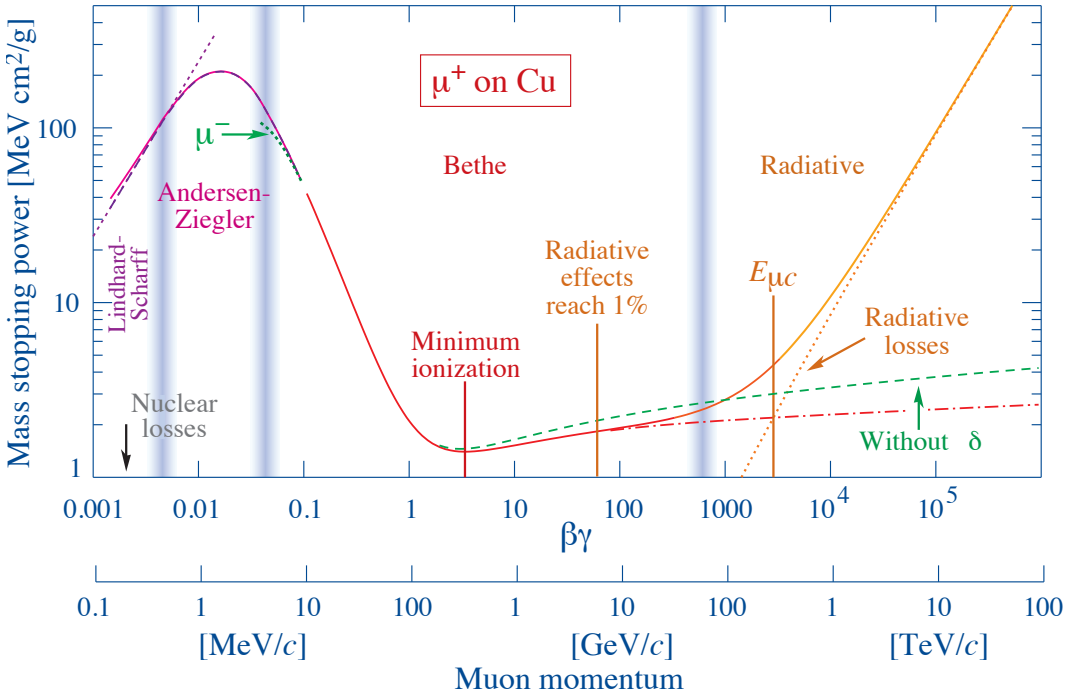

Figure 5: Mass stopping power (dE/dx) for positive muons in copper as a function of $\beta\gamma = p/Mc$ over nine orders of magnitude in momentum (12 orders of magnitude in kinetic energy). Solid curves indicate the total stopping power. Vertical bands indicate boundaries between different approximations. Source: Passage of Particles Through Matter from the Particle Data Group [17].

survive recapture through the recombination factor R, the deposited energy $\frac{dE}{dx}$ can be related to the ionized charge through:

$$\frac{dE}{dx} = \frac{W_{ion}}{R}\frac{dQ}{dx} \tag{4.5}$$

where $W_{ion} = 23.6\,\text{eV}$ represents the energy required to ionize an argon atom. Under the columnar model, recombination includes a significant bulk effect from the roughly cylindrical distribution of Ar^+ ions left in the trail of an ionizing particle [93]. The box model assumes a constant charge density from Ar^+ ions within a box region, and predicts the recombination factor as:

$$R_{Box} = \frac{1}{\beta(dE/dx)}\ln(\alpha + \beta\frac{dE}{dx}) \tag{4.6}$$

where $\alpha = 1$ is theoretically motivated and β is fit as a free parameter. However, given the assumptions made in this model, it fails to accurately describe reality at low dE/dx, so a modified box model is used where α is fit to 0.93 and β is fit to 0.3 cm/MeV.

4.3.3 Diffusion

LArTPCs detect ionized electrons in the detector volume after it has drifted to the anode plane. Ideally, this drift occurs in a uniform manner, with a constant drift velocity $v_d \approx 1\,\mathrm{m/ms}$ perpendicular to the anode plane. Then, simply multiplying the drift time t by the drift velocity would determine the location along the x-axis where the charge originated. However, in practice there are two significant real-world effects that must be considered for accurate charge reconstruction. First is the effect of diffusion, where the charge spreads out in a cloud over time from the pseudo-random walk of each particle. Because of the applied electric field, however, there is a preferred direction for electrons to drift, and so diffusion is parameterized by separate constants D_L and D_T for the longitudinal and transverse motion respectively. Through the electron mobility μ, the diffusion parameters can be computed from an extension of the Einstein-Smoluchowski relation [94]:

$$\begin{aligned} D_T &= \frac{kT}{e}\mu \\ D_L &= \frac{kT}{e}\left(\mu + E\frac{\partial\mu}{\partial E}\right) \end{aligned} \tag{4.7}$$

where kT is the electron temperature and e is the electric charge. The diffusion parameters control the rate at which an electron cloud grows, smearing the resolution from transverse diffusion and widening the signal pulse from longitudinal diffusion:

$$\sigma_t^2(t) = \sigma_t^2(0) + \left(\frac{2D_L}{v_d^2}\right)t \tag{4.8}$$

The passage of an electron past each wire plane motivates the non-zero width at $t = 0$ of $\sigma_t^2(0) \approx 1.6\,\mu$s.

4.3.4 The Space Charge Effect

While the discussion of ionization charge so far has focused on electrons, as they are the particles detected at the anode, it is also important to consider the impact caused by Ar^+ residing in the detector. Eventually Ar^+ ions will drift to the cathode and leave the fiducial volume; however, because of their much larger mass, they drift much slower and remain in the detector for an appreciable amount of time. Because the MicroBooNE detector is located on the Earth's surface and observes a significant flux of ionizing cosmic radiation, Ar^+ ions are constantly being generated throughout the detector, not just from neutrino interactions. The steady state background resulting from the roughly uniform production of Ar^+ ions and constant drift velocity towards the cathode is a linearly increasing charge density as you approach the cathode. This steady state is called space charge, and the impact that it has on electron drift through distorting the electric field lines in the detector is called the Space Charge Effect (SCE).

Since the positive space charge attracts negatively charged electrons, the SCE is seen primarily as pulling tracks towards the center of the detector. This creates an effective detector boundary, where charge ionized on the actual detector boundary will pass through the wire planes and be reconstructed inward of this boundary. Since the space charge distortions build over time (and therefore over drift distance), the effective detector boundary is largely unchanged near the anode, and becomes increasingly distorted towards the cathode. The effect was measured empirically by studying cosmic ray muons that cross the detector boundary, ionizing charge up to the true detector edge. The observed cutoff of these tracks, measured over a large sample, generates a reliable mapping from the effective to true detector boundary [95], as seen in Fig. 6. This mapping is especially important for dedicated cosmic

ray identification algorithms that leverage precision knowledge of the detector boundary.

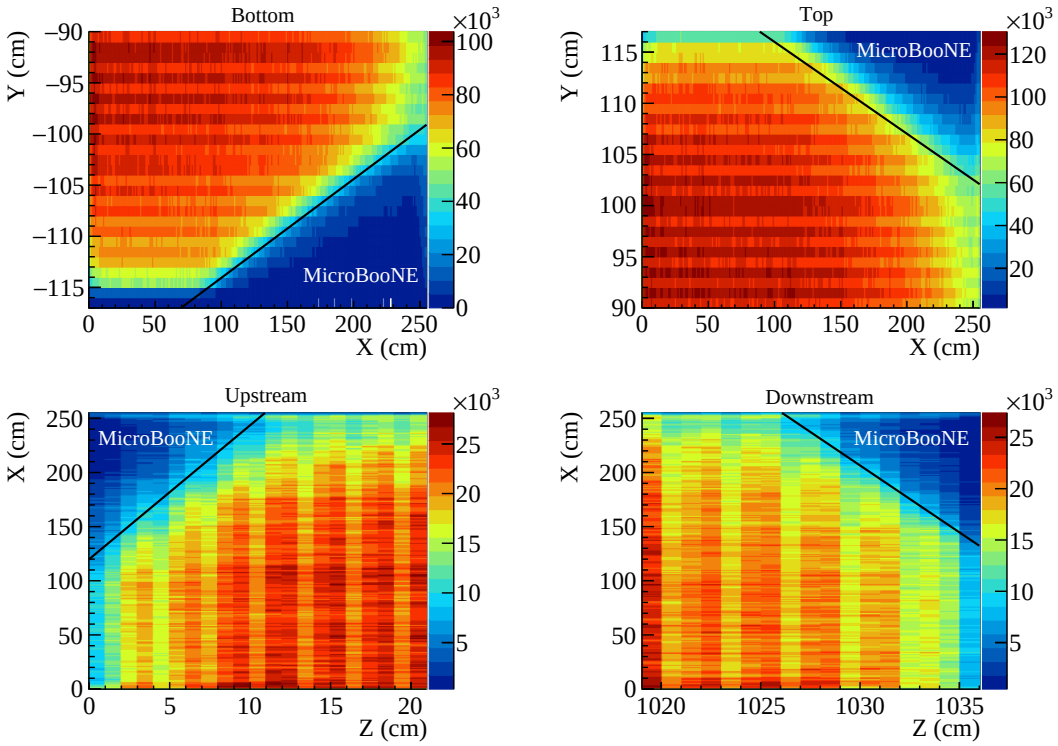

Figure 6: The effective detector boundary (black lines) at the four corners on the cathode side of the detector. The color scale shows the density of cosmic-muon charge clusters in the detector. Source: [95]

4.3.5 Charge Measurement

MicroBooNE has three wire planes located at the anode that serve as its primary particle detection instruments. The wire planes are flush with each other, with a 3 mm spacing between planes as well as a 3 mm pitch between wires on each plane [96]. When viewed along the drift dimension, the wires of different planes intersect at 60° angles, forming equilateral triangles. The first two planes in the path of drift electrons, labeled "U" and "V" respectively, each consist of 2400 wires at ±60° angles with the vertical, while the final "W" plane (sometimes called the "Y" plane) consists of 3456 wires oriented vertically.

The "U" and "V" planes are held at −110 V and 0 V respectively, and measure current

induced by the passage of electrons past them, earning them the name of induction planes. The "W" plane has a voltage of 230 V, drawing electrons onto it and giving it the label of collection plane. The signal measured on induction planes is bipolar as a result of electrons passing by them, while the signal on the collection plane is uni-polar. Since bipolar signals with a slight time offset interfere destructively, the collection plane is capable of producing more reliable charge measurements. Additionally, the "U" plane is exposed to variations in the voltage at the cathode while shielding the "V" and "W" planes behind it, making it the least reliable of the three planes.

The MicroBooNE detector uses a cold electronics system to amplify the current signal measured on the wires [96]. This is achieved using Application Specific Integrated Circuits (ASICs), with a gain of 14 mV/fC. By operating within close proximity to the wire planes and at liquid argon temperatures, the noise experienced in the electronics system is significantly reduced. This is crucial for enabling a low charge detection threshold, and is a significant advancement to the field of LArTPC design pioneered by MicroBooNE. After the signal is amplified by the cold electronics, it is sent outside the cryostat and digitized using a Data Acquisition system (DAQ). Current measurements are sampled at 2 MHz by the Analog-to-Digital Converter (ADC), creating a time tick interval of 0.5 μs.

Through de-convolution of the current measured on the wires of a plane, the time and 1D location of passage of electrons along the wire plane (perpendicular to the wire length) is determined, as shown in Fig. 7. By combining this information across two or more planes, the electron's 2D location of passage on the wire plane can be determined. In the case of MicroBooNE, the use of three wire planes gives a level of redundancy for a more accurate position measurement. Full 3D positional information is later achieved by determining the drift time from the time of passage and the time of interaction as measured through the light system, and then multiplying by the drift velocity.

For MicroBooNE's electric field of 273.9 V/cm, the electron drift velocity is approximately

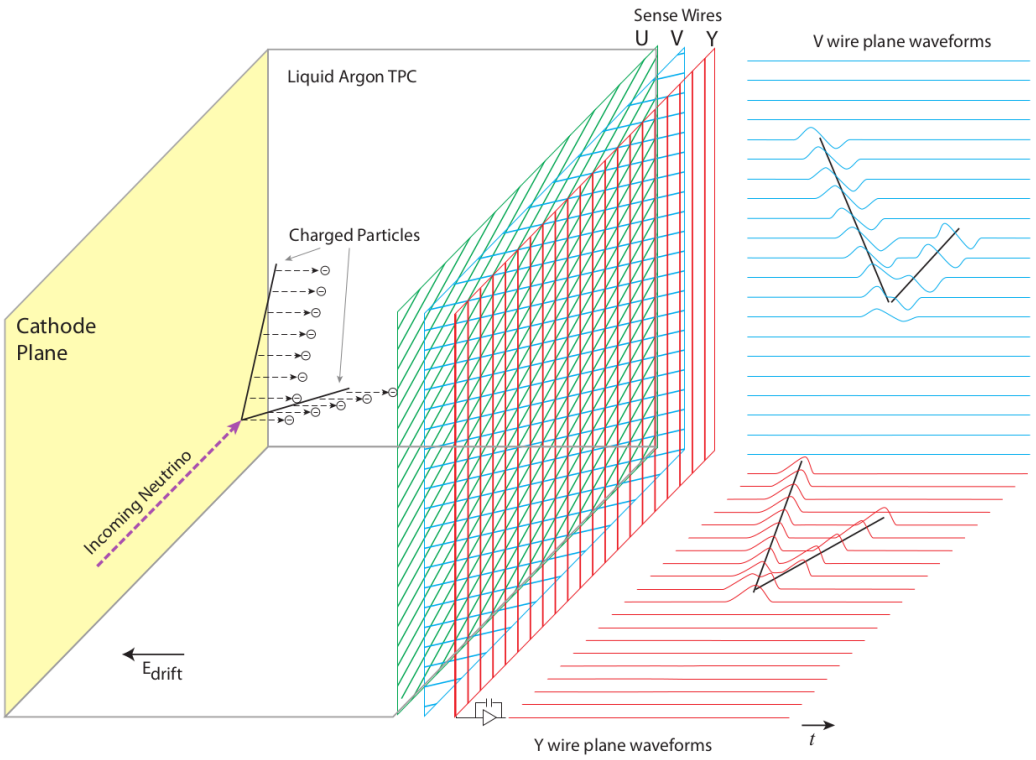

Figure 7: Diagram showing how charge is measured using the wire planes in a TPC.

1 m/ms, with a total drift time from cathode to anode of 2.3 ms, called the drift window. The full readout window is 4.8 ms long, comprised of three 1.6 ms windows corresponding to the expected drift window time at a nominal voltage of 500 V/cm (which MicroBooNE was unable to fully achieve). The readout window begins one 1.6 ms window before the beam spill and continues a further two nominal drift windows to help with cosmic ray detection. Any charge measured within one drift window of a BNB beam spill could potentially be the result of a neutrino interaction, allowing for a significant cosmic ray background as a result of the inherently slow drift velocity of electrons. On average, there are 20-30 cosmic rays in each drift window, with only one neutrino per ~ 600 beam spills. Reducing this source of background is the primary goal of the PMT light system, which measures the prompt scintillation light released by Ar following a neutrino interaction.

4.3.6 Light Collection System

The passage of ionizing particles through argon can create two excited states, ${}^1\Sigma_\mu^+$ and ${}^3\Sigma_\mu^+$ [92]. These decay with half-lives of $\tau_s = 5\,\text{ns}$ and $\tau_3 = 2.1\,\mu\text{s}$, emitting a photon at 128 nm that LAr is optically transparent to. As a result, ionizing particles created in a neutrino interaction create a signature of prompt Ar scintillation light with two decay widths that is easily detected.

An array of 32 PMTs is arranged behind the anode plane to capture this scintillation light, with a digitized time resolution of 15.6 ns. Each PMT first has a plate coated in tetraphenyl-butadiene, which absorbs the 128 nm scintillation light and re-emits light at $\sim 425\,\text{nm}$, which the PMT photocathode is sensitive to. Through the photoelectric effect, photons striking the photocathode eject electrons [97]. These cascade toward the anode, increasing the strength of the signal measured, given in units of Photo-Electrons (PEs). The precision timing of the PMT system allows for timing of ionizing particles to be determined to a much higher resolution than either the TPC drift window (2.3 ms) or the BNB beam spill ($1.6\,\mu\text{s}$). After matching a PMT light flash to a TPC charge cluster, its 3D position and time of origin can be well determined for use in event selection.

The PMT array collects data over a $23.4\,\mu\text{s}$ window for each neutrino beam spill reported by the Fermilab accelerator division. Combined with the 4.8 ms TPC readout window, this selective data taking in sync with beam spills is called the hardware trigger. However, the hardware trigger alone is not selective enough, as only 1 in 600 beam spills contain a neutrino interaction. A software trigger is also applied, which rejects data taken from any readout window that contains less than 9.5 PEs of energy, as measured by the PMT array. Together, the hardware and software triggers greatly reduces cosmic ray background contamination in event selections, and greatly reduce the amount of information storage required.

5 Wire-Cell Event Reconstruction

The MicroBooNE collaboration has developed three separate analysis chains to reconstruct and select neutrino events inside the detector. This analysis is based on the Wire-Cell reconstruction, which takes a philosophy of preserving as much information as possible through the reconstruction chain. Sec. 5.1 describes the process of computing charge measurements from the current readout on wires through a 2D deconvolution algorithm that leverages the correlated info across wires. Sec. 5.2 describes the process of combining charge measurements across wire planes to form a fully 3D representation of the measured charge, followed by steps that cluster this charge into fully-connected interactions. The 3D charge clusters are used in a global fit of charge clusters to light flashes, described in Sec. 5.3. Then, Sec. 5.4 describes how these charge clusters are parsed through the use of a graph-based framework that allows for accurate determination of particle trajectories and energy deposition rates. Together, these algorithms serialize the information-dense raw data produced in the LArTPC into a high-level description that contains a vast amount of the salient features of the data, leading to the overall particle flow diagram and neutrino vertex identification in Sec. 5.5. This forms the foundation on which the selection algorithms in Sec. 6 rely to produce a high-quality neutrino selection. The algorithms described in this chapter each represent advancements in the field of LArTPC reconstruction tools, and together comprise the Wire-Cell reconstruction chain, which has pioneered a maximally-information-preserving technique of event reconstruction.

5.1 Signal Processing

Signal Processing covers the steps involved in taking the raw current measurements on the LArTPC wires and producing a measurement of the charge that passed by each wire. This is roughly split into two parts: the modeling of the electric field response in the presence of moving charge, described in Sec. 5.1.1, followed by the extraction of the charge distribution

from the convolution of the current measurement with the detector response, described in Sec. 5.1.2. Wire-Cell builds upon existing techniques by developing a 2D deconvolution algorithm that uses the combined current measurements across nearby wires to more accurately determine the charge signal at each wire. This is a substantial advancement to the field, and the technique has since been adopted by all MicroBooNE analyses. As the beginning of the reconstruction chain, improved extraction of the charge distribution allows for more accurate particle reconstruction at all stages downstream.

5.1.1 Field Response

To convert the measured current signals on each wire to accurate charge estimates, it is important to precisely describe the field response generated on each wire by drifting electrons. This relation can be found by considering the electric field along the length of a wire, as produced by an electron and across the electron's drift path. However, it is simpler to use newton's 3rd law to instead consider the electric field produced by a wire, $\vec{E}(\vec{r})$. This is achieved by simulating the field response in the MicroBooNE detector with the Garfield [98] software package. A 2D simulation was used, shown schematically in Fig. 8, in which a range of 10 wires on either side of the wire of interest were used for each of the three wire planes. To compute the electron drift path, the electron was simulated with a starting position 10 cm ahead of the "U" plane, across a range of transverse starting locations, each 0.3 mm apart. The simulated drift path and electric potential at each location are shown in Fig.9.

Electron drift $\vec{v}_d$ in LArTPCs is slow enough to allow for a quasi-static approximation, under which Ramo's theorem [100] describes the current $i(\vec{r})$ induced by an electron cloud of charge q as a function of the drift velocity and the electric field of a wire at unit potential:

$$i(\vec{r}) = -q\vec{E}(\vec{r}) \cdot \vec{v_d} \tag{5.1}$$

From combining Eqn. 5.1 with Fig. 9, a few key insights can be made. The drift paths and

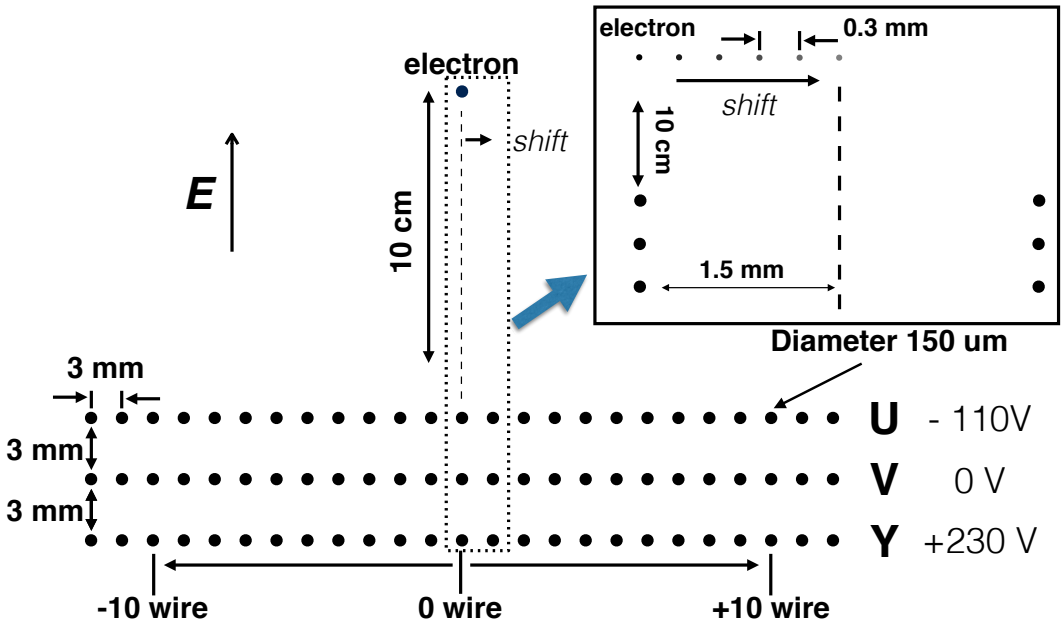

Figure 8: Illustration of the 2D Garfield simulation scheme (dimensions not to scale), where black dots indicate individual wires. MicroBooNE's anode plane-to-plane spacing is 3 mm, with 3 mm wire pitch in each plane. The inset denotes the sub-pitch designation of electron drift paths whereupon the field response is calculated. Source: [99]

potential maps suggest a significant difference in current induced between charge incident on a wire, and charge originating at a location equidistant between two wires. This difference is shown in Fig. 10. Moreover, there is significant current induced on the wires in the local area beyond the closest wire. These effects are plane dependent, with the un-shielded "U" plane experiencing the largest impact from distant charge, both in the transverse and longitudinal directions.

These effects combine to create significant differences in wire response based on the larger topology of the ionized charge. Shifting focus beyond a single point of drift charge, the full 3D structure of a cluster of charge can impact how overall current signals are measured, and how charge can be reconstructed. For simplicity (and largely adhering to the reality of particle trajectories), a linear charge distribution is considered, which is defined by two angles. θ_{xz} describes the direction within the xz plane, with 0° corresponding to the beam direction z and 90° corresponding to the drift direction x, while θ_y is the angle off of the wire

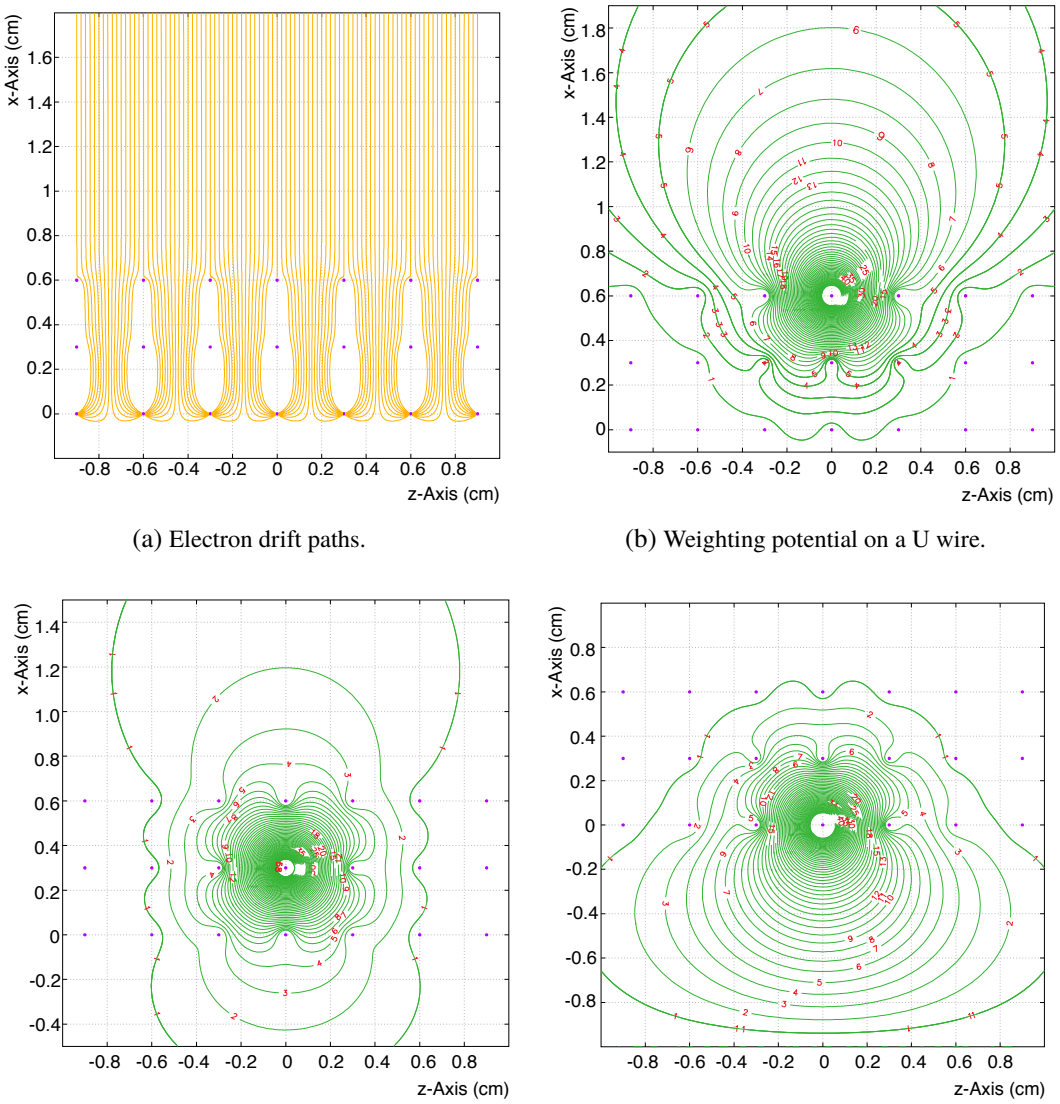

(a) Electron drift paths.

(b) Weighting potential on a U wire.

(c) Weighting potential on a V wire.

(d) Weighting potential on a Y wire.

Figure 9: A demonstration of electron drift paths in the applied electric field (panel a) and weighting potentials (panels b, c, d) on individual wires of the 2D MicroBooNE TPC model, using the Garfield program. The x-Axis is in the drift field direction and the z-Axis is in the beam direction. Values for the weighting potential are indicated in percentage on each equipotential line, ranging from 1% for the farthest to 60% for the closest illustrated. Source: [99]

direction y (not to be confused with the overall detector vertical dimension y).

The current induced from a 1 m MIP track is shown in Fig. 11 for various configurations of θ_{xz} and θ_{xz}. There are three extreme cases to consider. First is the isochronous track, considered in Fig. 11b, where $\theta_{xz} = 0$ and all ionized charge arrives at the wire planes simultaneously.

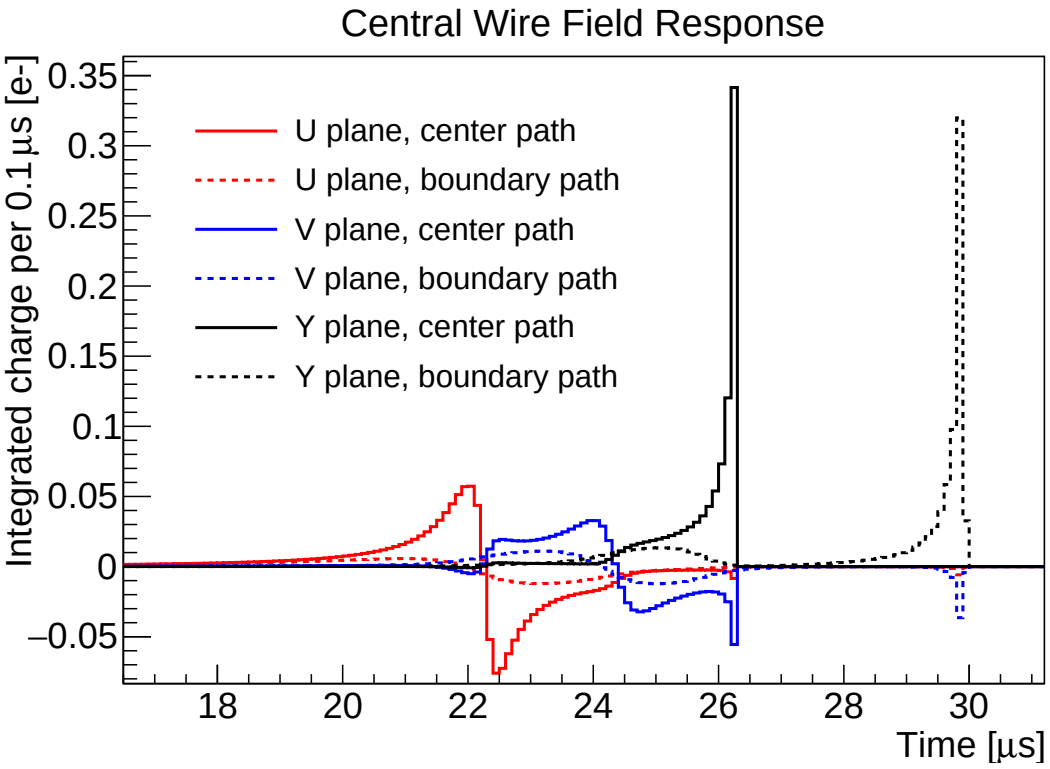

Figure 10: Induced current from a center path (incident on sensing wire) and boundary path (1.5 mm transverse offset, equidistant between wires) of one drifting ionization electron for the three wire planes. The Y-axis is the integrated charge over 0.1 µs intervals. Source: [99]

In this case, long-range electric field effects are important across all wires, especially those in the induction plane. However, time-based cancellation effects on bipolar signals do not play a role. The next case is the prolonged track, where $\theta_{xz} = 90°$. In this situation, there is maximal destructive interference in the induction plane wire signals resulting from charge at different time offsets. As a result, it can be particularly difficult to accurately measure charge on the induction planes.

Finally, there are aligned tracks, where a track is (nearly) parallel with the wires of a particular plane, although of course this can never happen to multiple planes simultaneously. In this case, the aligned plane is useless at determining the distribution of charge along the particle trajectory, leaving the job to the other two planes. This can make successful trajectory fitting in Sec. 5.4 difficult, and can be particularly problematic when the aligned wire plane is

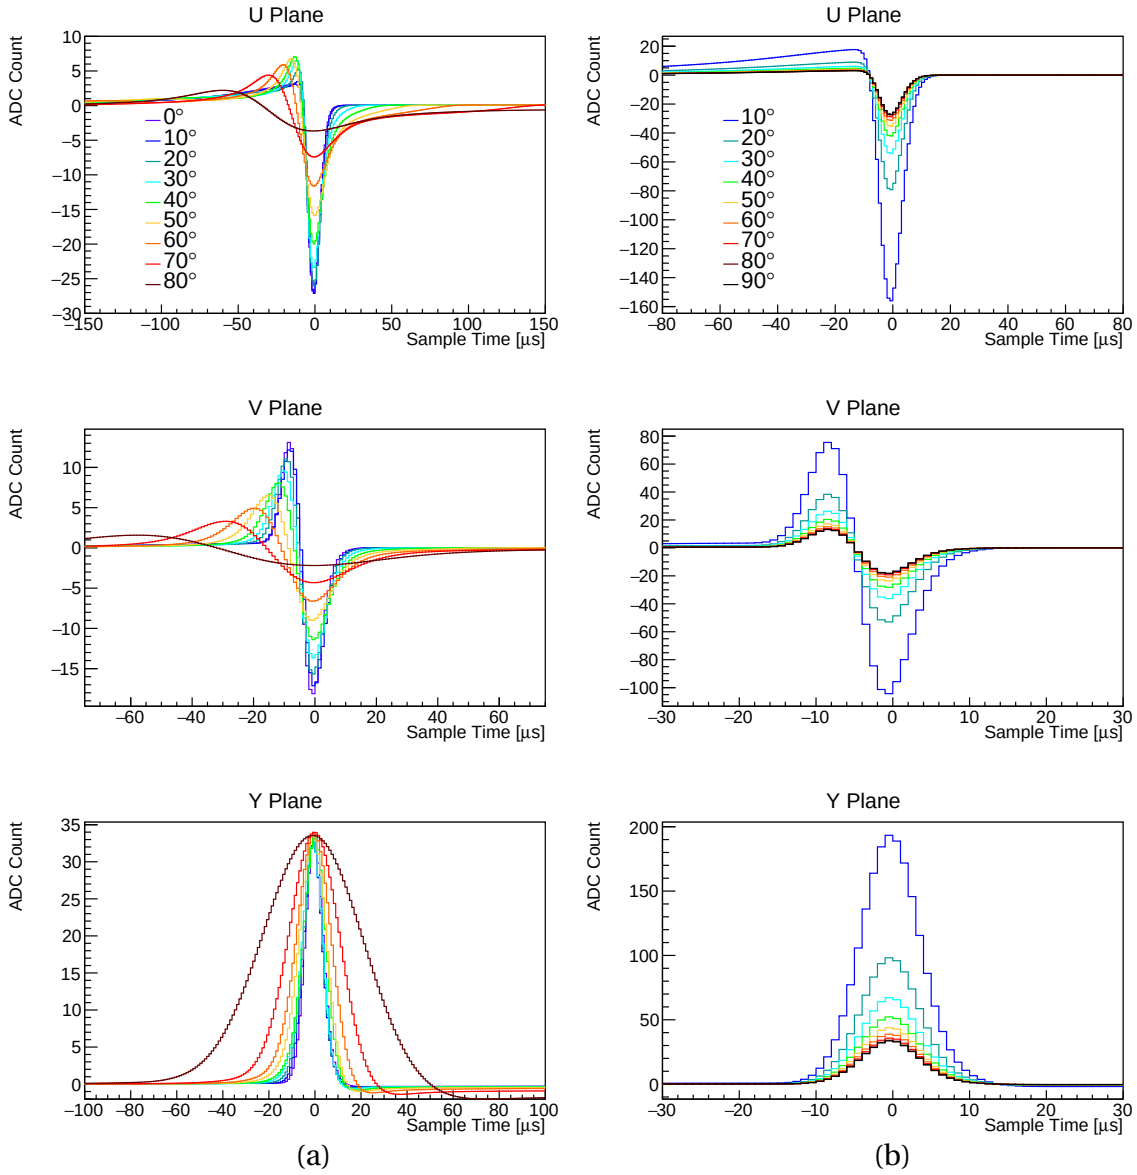

Figure 11: Simulated baseline-subtracted MicroBooNE TPC signals for a 1 meter long MIP track (~ 4400 ionization electrons per mm). Detector physics effects and the nominal MicroBooNE electronics response [101] were included. Left: the track is traveling perpendicular to each wire plane orientation ($\theta_y = 90°$) with θ_{xz} varying in the xz plane with respect to the z-axis. Right: the track is isochronous ($\theta_{xz} = 0$) with θ_y varying with respect to the wire orientation. Source: [99]

the collection plane and the induction planes' signals suffer from destructive interference or high noise (such as from the cathode), or when one or more of the required planes suffer from dead wires in the region. These difficulties will be discussed in more detail in Sec. 5.4.

5.1.2 Charge Measurement

With the induced current across wires accurately modeled, it is time to compute the underlying charge distribution for a given current measurement on each wire. The current measurement $M(t')$ can be seen as a convolution of the desired charge signal distribution $S(t)$ and the detector response $R(t,t')$, which maps the current measurement at some time t' to an element of original charge signal at time t:

$$\begin{aligned} M(t') &= \int_{-\infty}^{\infty} R(t,t') \cdot S(t) dt \\ &= \int_{-\infty}^{\infty} R(t'-t) \cdot S(t) dt \end{aligned} \tag{5.2}$$

where the time invariance of the response function has been asserted in the second line. Inverting this relation to solve for $S(t)$ involves the method of deconvolution, and has been performed in 1D in previous LArTPC experiments [102]. This is achieved by taking the Fourier transform to express the relation across the angular frequency ω, and then solving for $S(\omega)$:

$$S(\omega) = \frac{M(\omega)}{R(\omega)} \tag{5.3}$$

The charge signal distribution over the time domain $S(t)$ can be obtained by simply applying the inverse Fourier transform to $S(\omega)$.

This straightforward procedure encounters a couple of difficulties from real-world considerations. First, the measured current contains noise from the electronics [103], and the mathematical solution $R(\omega)$ to the Fourier transform ignores practical limitations of detector equipment, by which $R(\omega) \to 0$ at large ω. This necessitates the introduction of a filter function $F(\omega)$ to modify the effective detector response description by attenuating high-frequency

noise and respecting the limitations of a realistic detector:

$$S(\omega) = \frac{M(\omega)}{R(\omega)} \cdot F(\omega) \tag{5.4}$$

The filter function used is derived from the Wiener filter [104], which suppresses regions of high expected noise $\overline{N^2(\omega)}$ relative to signal strength $\overline{S^2(\omega)}$:

$$F_W(\omega) = \frac{\overline{R^2(\omega)S^2(\omega)}}{\overline{R^2(\omega)S^2(\omega)} + \overline{N^2(\omega)}} \tag{5.5}$$

The Wiener filter has the advantage of maximizing the signal-to-noise ratio with minimum mean-squared-error of the deconvolved distribution. However, the Wiener filter does not conserve the overall charge signal measurement ($\lim_{\omega\to 0} F_W(\omega) < 1$), and smears the distribution non-locally over the time domain (and therefore the drift dimension). To prevent these issues, a Wiener-inspired filter is used:

$$F(\omega) = \begin{cases} e^{-\frac{1}{2}\left(\frac{\omega}{a}\right)^b} & \omega > 0 \\ 0 & \omega = 0 \end{cases} \tag{5.6}$$

where a and b are free parameters, and the function $F(\omega)$ is fit to the Wiener filter $F_W(\omega)$. By setting $F(\omega = 0) = 0$, the filter removes any DC component of the current measurement, while $\lim_{\omega\to 0} F_W(\omega) = 1$ ensures that the overall charge normalization is not affected. The decaying Gaussian tail suppresses non-local smearing effects on the distribution of charge.

The procedure as presented so far is mathematically valid for extracting the signal charge from the current measurement on a wire-by-wire basis. However, it does not include the current information from nearby wires, which has been shown to be significant. Under real-world conditions, with noise and measurement uncertainty, as well as destructive interference on induction planes, the additional information measured on nearby wires is particularly helpful for an accurate charge measurement.

The information from nearby wires can be incorporated by redefining the measured current $M_i(t')$ on wire number i at time t' as a sum of charge distributions $S_{i+j}(t)$ associated with their closest wire, over the $2k+1$ local wires $j \in \{-k, ..., k\}$ (as well as the convolution over time domain):

$$M_i(t') = \int_{-\infty}^{\infty} \sum_{j=-k}^{k} R_j(t'-t) \cdot S_{i+j}(t)\, dt \tag{5.7}$$

This essentially represents the measured current as a 2D convolution, over both time and wire number. As before, applying a Fourier transform solves the convolution in time:

$$M_i(\omega) = \sum_{j=-k}^{k} R_j(\omega) \cdot S_{i+j}(\omega) \tag{5.8}$$

What is left is a system of linear equations, and can be simplified by asserting that the response function should be symmetric in wire number ($R_{-j}(\omega) = R_j(\omega)$):

$$\begin{pmatrix} M_1(\omega) \\ M_2(\omega) \\ \vdots \\ M_{n-1}(\omega) \\ M_n(\omega) \end{pmatrix} = \begin{pmatrix} R_0(\omega) & R_1(\omega) & R_{n-2}(\omega) & R_{n-1}(\omega) \\ R_1(\omega) & R_0(\omega) & R_{n-3}(\omega) & R_{n-2}(\omega) \\ \vdots & \vdots & \ddots & \vdots \\ R_{n-2}(\omega) & R_{n-3}(\omega) & R_0(\omega) & R_1(\omega) \\ R_{n-1}(\omega) & R_{n-2}(\omega) & R_1(\omega) & R_0(\omega) \end{pmatrix} \begin{pmatrix} S_1(\omega) \\ S_2(\omega) \\ \vdots \\ S_{n-1}(\omega) \\ S_n(\omega) \end{pmatrix} \tag{5.9}$$

where $R_j(\omega)$ should be 0 for any $j > k$ corresponding to the response function for charge outside the local region of a wire.

Since the matrix R is symmetric and Toeplitz (meaning that each diagonal descending from left to right has all elements equal), the matrix inversion can be computed using a discrete Fourier transform [105]. Fig. 12 shows the resulting charge signal extracted before and after both 1D and 2D deconvolution, highlighting the sensitivity to signal topology and the robust performance of the 2D deconvolution.

In theory, the deconvolution in time could be performed over the entire readout window.

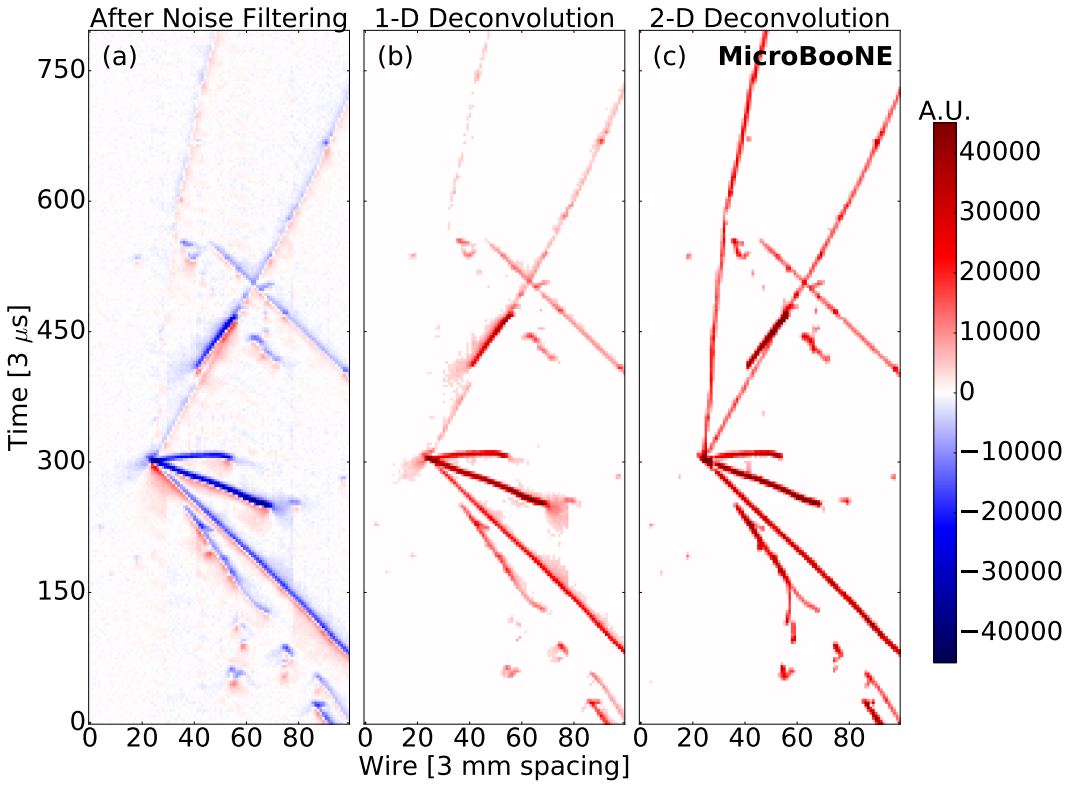

Figure 12: A neutrino candidate from MicroBooNE data (event 41075, run 3493) measured on the "U" plane. (a) Raw waveform after noise filtering in units of average baseline subtracted ADC scaled by 250 per 3 us. (b) Charge spectrum in units of electrons per 3 us after signal processing with 1D deconvolution. (c) Charge spectrum in units of electrons per 3 us after signal processing with 2D deconvolution. Source: [99]

However, this would be a slow process, and more importantly leave the deconvolution vulnerable to low-frequency noise. The wider a time window used, the longer the period of noise allowed through a Fourier transform. In theory, a high-pass filter could be applied to suppress this noise similar to how the Wiener filter suppresses high-frequency noise, however this could destroy real charge information, as well as smear the calculated charge distribution. Instead, a Region of Interest (ROI) is used, which is a time interval surrounding each period of current measurement on a wire. The deconvolution is only performed within this time window, avoiding sources of noise outside the ROI and corresponding to low-frequency noise.

5.2 Tomographic Imaging

The goal of imaging is to tomographically reconstruct a fully 3D distribution of charge from the charge measurements across each wire plane as a function of time. Moreover, this reconstruction should be topology independent and preserve as much of the physical information captured in the wire measurements, so as to allow for the widest range of accurate particle reconstruction and event selection downstream of imaging. After a 3D distribution is constructed, it is useful to group charge into connected components called clusters, ultimately so that a single interaction is fully grouped as one cluster. This allows downstream algorithms to be written without having to significantly consider whether the scope of charge considered is too wide or too narrow, allowing for efficient high-level feature extraction.

These goals are made difficult by some of the limitations in LArTPC capabilities, as well as imperfections of the MicroBooNE detector. Typically, tomography is performed by combining many lower dimensional projections to reproduce the original image, such as in medical applications including MRIs and CT scans. However, in the case of LArTPCs, only three projections corresponding to the three wire planes are available. This makes reconstructing an image more difficult, and in general even with perfect information no unique solution mapping from wires to 3D charge exists. It is also possible to reconstruct so-called ghosts, where a coincidental intersection of wires with simultaneously measured charge suggests a solution that is in fact non-physical. These issues are greatly exacerbated by the existence of dead wires within MicroBooNE, which span $\sim 10\%$ of each wire plane, covering $\sim 30\%$ of the wire plane surface area. In these regions, only two wire planes (or fewer!) are functional, further reducing the information available to perform tomography with.

To maintain fidelity in the reconstruction despite these obstacles, the full range of detector information is used, as well as a number of key insights connecting wire measurements and characterizing the structure of the 3D imaging solution. The wire measurements provide information on the time of arrival, 1D location along wire number for each plane, and

magnitude of charge observed. Additionally, since induction planes are set at a potential so as to not collect charge, we can apply charge conservation between wire plane measurements to aid in the reconstruction. Furthermore, charge drifting towards the anode is always the same polarity and never reverses direction, implying a condition of positivity in the solution. Finally, the physics of ionizing particles motivates the use of a proximity constraint, by which continuous charge distributions, such as particle tracks, are preferred.

The imaging detailed in this section is able to reconstruct a fully 3D charge distribution with non-functional regions totalling only 3% of the wire plane area. This is achieved through a robust reconstruction algorithm that is capable of reconstructing charge in regions with only two functional wire planes, aided by more precise 3-plane measurements in the surrounding areas. The 3D charge distribution reconstruction is particularly noteworthy, as it enables a suite of high-level particle reconstruction algorithms, some of which are discussed in Sec. 5.4. The completion of 3D charge reconstruction represents turning a corner from unraveling the convoluted measurements produced by TPC wires towards freely using detector information as presented in its most physically fundamental form to identify particle types, energies, and other kinematics.

5.2.1 Tiling

Before solving for a 3D charge distribution, wire info needs to be appropriately bundled to reduce the computational time needed and complexity encountered by a solving algorithm. Furthermore, as a result of smearing from diffusion, noise filtering, and residual long-range charge-current induction, conservation of charge across planes is more accurate when applied over a larger area. Obtaining a consistent charge measurement across wire planes is critical to solving the imaging equation 5.10 in Sec. 5.2.2, therefore bundling together wires corresponding to a region of charge is necessary [106]. Additionally, tiling serves as a conservative estimate on the possible 3D locations of charge, limiting the scope of later

charge solving and and deghosting steps to one of deciding where charge is not.

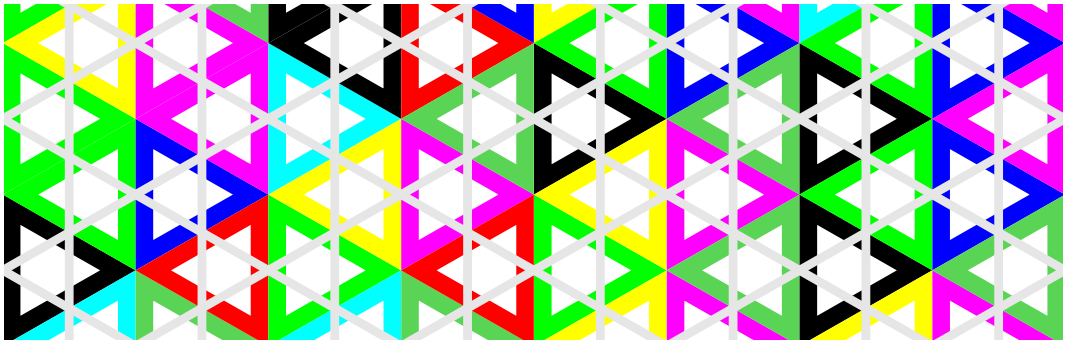

Figure 13: Cells constructed with the MicroBooNE detector geometry, colored differently for visibility. Wires are colored in grey. Each cell is an equilateral triangle as a result of the detector geometry of ±60° angles between wires across planes. Source: [107]

First wire measurements over time are merged into 2 μs time slices consisting of four sampling points of TPC readout. This introduces negligible information loss because noise filtering heavily suppresses signals above 0.25 MHz. Next fired wires can be bundled based on their proximity. To do this, the concept of a “cell" is defined, as a primitive geometric object that maps reconstructed charge from its nearest associated wires to a 2D position. For three wires on different planes that nearly intersect, they will bound a smallest-size equilateral triangle across the planes. The center of each of such triangle is used to seed a Voronoi cell [108], which together divide the plane, as shown in Fig. 13.

In the simplest case, three wires across planes that bound a cell and registered as “fired" (indicating the presence of nearby charge) would map to the bounded cell being labeled as “hit" [107]. Note that in general, many wires can fire and the association to hit cells is not so simple, such as in Fig. 14. Tiling is the process by which multiple adjacent hit cells are grouped into a blob, defined by the consecutive fired wires on each plane that span the hit. Fig. 15 shows a region of hit cells (with centers marked by blue dots) and corresponding fired wires (red) that are tiled together to form a single blob, outlined in blue. The boundary of a single cell is also shown, in black.

Note that this definition of a blob requires all three wire planes to be functioning in the

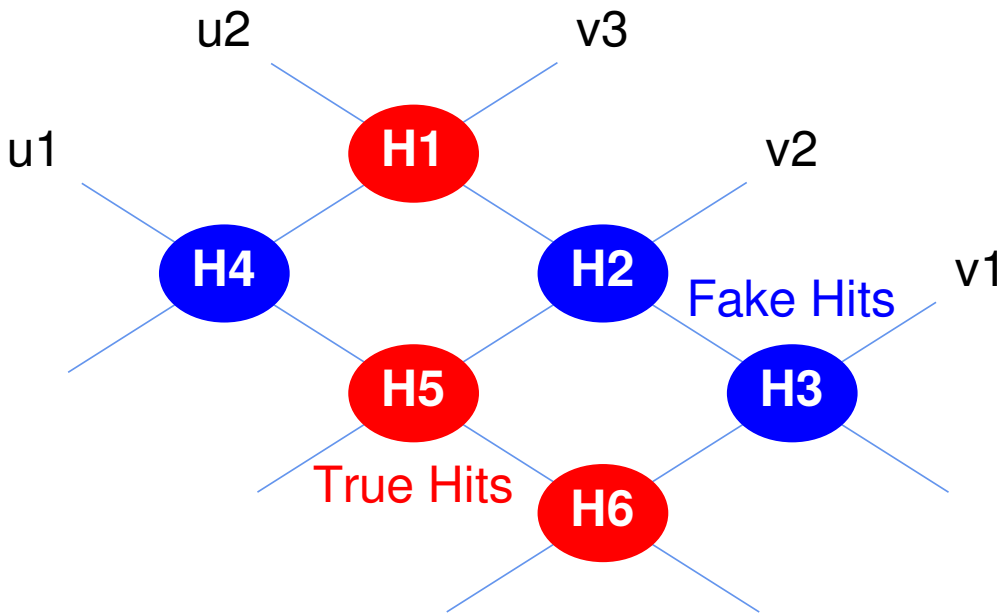

Figure 14: Illustration of the hit ambiguity caused by the wire readout in a simplified two-wire-plane example. This examples imagines sets of two and three wires that fire, creating six potential intersections, but only three correspond to true charge hits. Source: [107]

region. This requirement is at odds with the reality of the MicroBooNE detector, which suffers from ~ 10% dead wires on each wire plane. This means that ~ 30% of cells are bounded by at least one dead wire, shown visually in Fig. 16. This represents a significant loss in selection efficiency corresponding to events originating in a dead region. Even among surviving events, the chance of losing information across part of an interaction threatens the overall ability to correctly identify particles and estimate energies, harming both oscillation and cross section measurements.

Motivated by these dire prospects, the requirements to label a cell as hit are relaxed. In addition to cells bound by three fired wires, hit cells are also allowed to be bound by two fired and one non-functional wire. This redefinition reduces the non-functional regions from ~ 30% of the wire-plane area to ~ 3%, at the cost of greatly increasing the number of false hits, called ghosts. Although some ghosts are always possible, as demonstrated in Fig. 14, they are

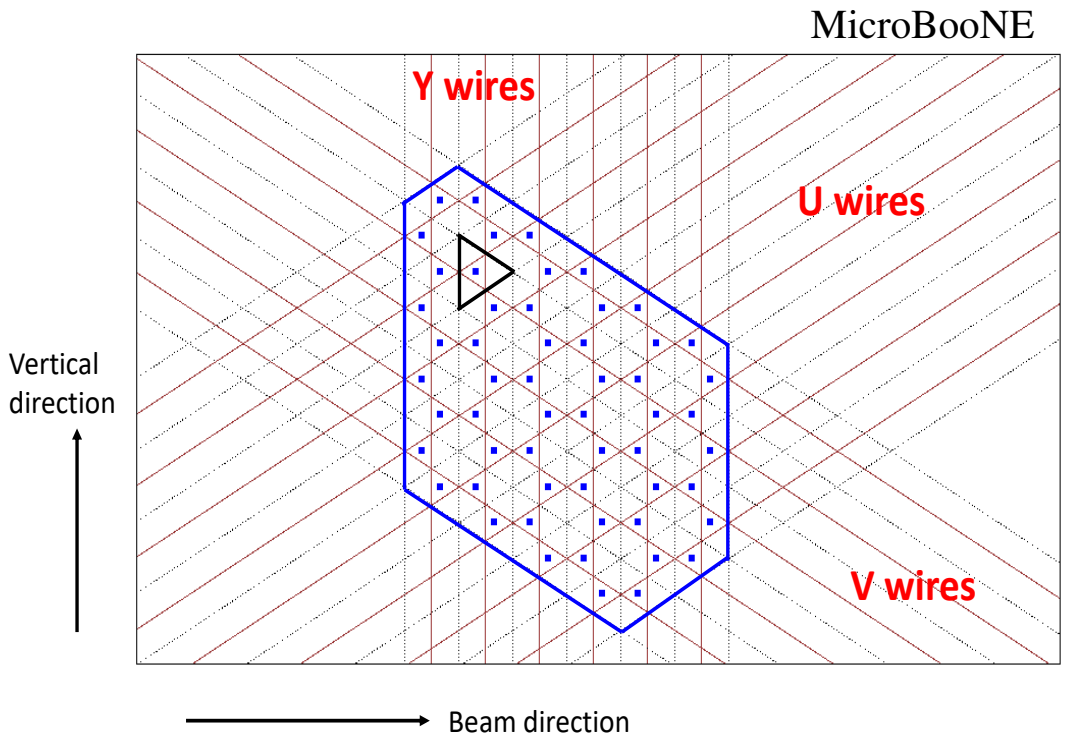

Figure 15: An example of the hit cells and blob constructed by fired wires. Each wire is represented by a solid red line and the wire boundaries are represented by dashed black lines. All hit cells have equilateral triangular shapes and are marked with blue dots at their centers. An example cell is marked by the black triangle. A blob is formed by the contiguous hit cells and marked by solid blue lines. Source: [106]

far more prevalent when ~ 10% of the wires are essentially treated as always fired. This issue will be addressed in detail in Sec. 5.2.3.

5.2.2 Charge Solving

The core goal of imaging in a LArTPC is solving the relation between the charge of fired wires and the 3D distribution of charge across time and cells. After tiling, this becomes a relation between merged wires and blobs, and can be expressed through a system of linear equations called the imaging equation:

$$y = Ax \tag{5.10}$$

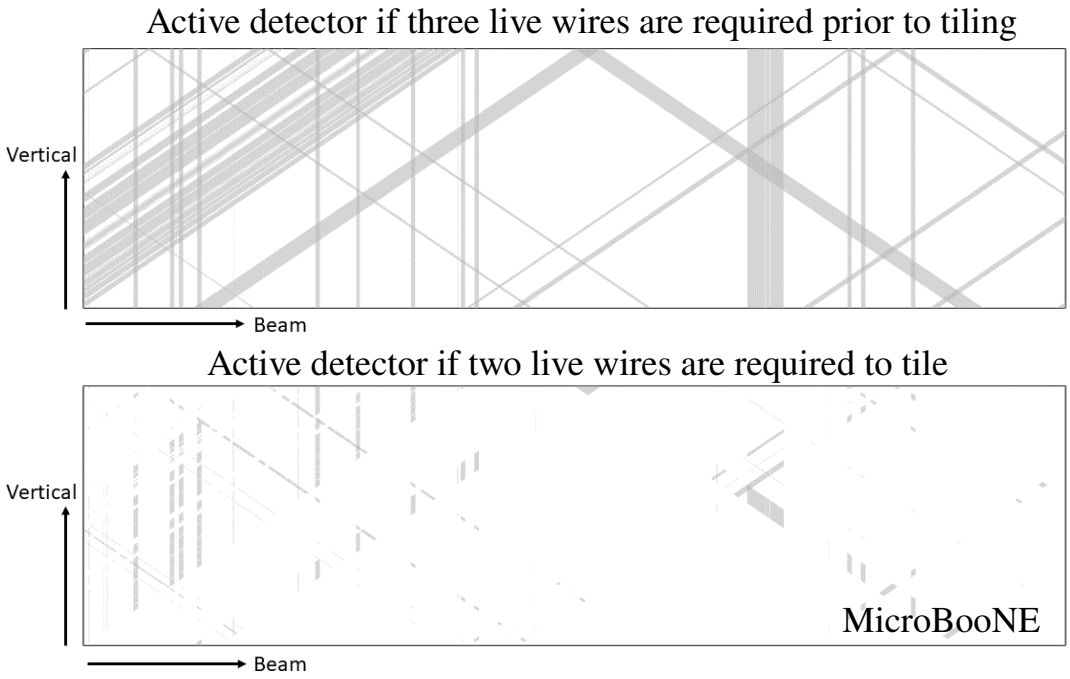

Figure 16: Visualization of non-functional detector regions in grey under a three-plane requirement (totaling ~ 30% non-functional) in the top panel and a two-plane requirement (totaling ~ 3% non-functional) in the bottom panel. Source: [106]

where y is a vector of charges across fired wires, x is a vector of charge across hit cells, and A is a matrix with elements $A_{ij} = 1$ if and only if the cell x_j has been grouped with wire y_i through tiling, and 0 otherwise. This system of equations is inherently under-determined [107], as there are $\sim n^2$ hit cells corresponding to the $\sim n$ fired wires on each plane; even in Fig. 15 there are 15 hit cells and only 11 fired wires. This prevents the existence of a unique solution, even before considering measurement uncertainty, noise filtering, and non-functional wires.

Instead of attempting to directly solve Eqn. 5.10, it is important to leverage our physical knowledge and intuition regarding the system. First, it is important to apply the positivity constraint. Only electrons drift toward the anode, and they never reverse direction, meaning that the charge solution can be described through a distribution of a single polarity (in this case, positive is used). Additionally, LArTPC activity is sparse; there are far more un-fired than fired wires at any time. When considering the distribution of charge across hit cells,

the true solution is most likely to be the one with the most ghost hits and correspondingly the fewest true hits with nonzero charge. This intuition can be applied mathematically by shifting from exactly solving a system of linear equations to minimizing a test statistic T_0 that measures the number of hits with nonzero charge, while obeying the relations in Eqn. 5.10. This is achieved through the mathematics of the ℓ_p norm, specifically the ℓ_0 norm in this case:

$$\ell_p(x) = ||x||_p = \left(\sum_i |x|^p \right)^{1/p} \tag{5.11}$$

$$T_0 = ||x||_0 \qquad \text{(subject to } y = Ax\text{)} \tag{5.12}$$

While minimizing T_0 is sufficient to find the best practical solution to the imaging equation, it is often not computationally feasible. It is essentially a combinatorics problem of testing all possible configurations, with an NP-hard complexity. This is a problem encountered elsewhere in the field of tomography [109], where the technique of compressed sensing [110, 111] allows for the approximation of the ℓ_0 norm solution with a much faster time complexity. This is achieved by replacing the ℓ_0 norm with the ℓ_1 norm, which well approximates the ℓ_0 solution and is convex in x with a global minimum, allowing for speedy minimization via gradient descent.

This test statistic can be improved by applying the physical knowledge of proximity; touching blobs over adjacent time slices are more likely physical, and so their ℓ_1 penalty term should be reduced. A series of weights ω is applied to the cells in x, scaling them down by a^n where n is the number of blobs connected in time to the target blob and a is a scaling factor. The test statistic is further modified by allowing some flexibility in the solution to Eqn. 5.10 to account for measurement uncertainty. The measurement and signal vectors y and x are normalized with the covariance matrix V of charge measurement uncertainties using Cholesky decomposition: $V^{-1} = Q^T Q$. The new test statistic T_1 to be minimized is

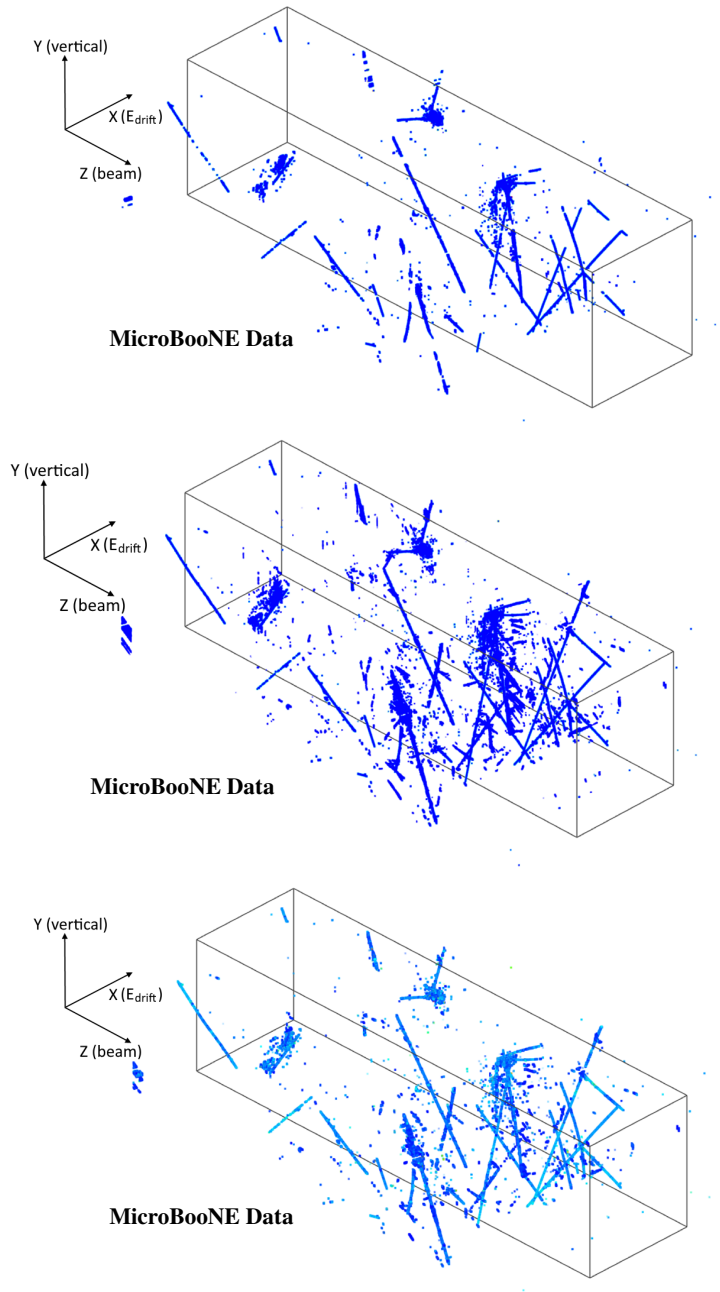

Figure 17: Comparison of the tiling results and the charge solving result from MicroBooNE data (event 41075, run 3493). The solid black box represents the LArTPC active volume with an x-position (converted from the readout time) relative to the neutrino interaction time. Top: tiling result under the strict 3-plane requirement. Middle: tiling result under the relaxed 2-plane requirement. Bottom: solved charge distribution from the blobs generated under the 2-plane requirement. The color scale represents the resulting charge values in the charge solving. Source: [106]

given by:

$$T_1 = ||Qy - QAx||_2^2 + \lambda||\omega \cdot x||_1 \tag{5.13}$$

where λ controls the strength of the ℓ_1 norm.

Fig. 17 demonstrates the results of charge solving as discussed in this section. The set of blobs under 3-plane and 2-plane requirements are contrasted with the solved distribution of charge, which removes many of the ghosts introduced in the 2-plane case. However, many ghosts still remain, motivating a dedicated effort to remove them through a repeated deghosting process.

5.2.3 Deghosting

The reason that the imaging equation misses so many ghosts is because it is only attempting to reduce positional redundancy, with just a slight modification in the form of the weights ω to connect the information across time slices. Therefore, the overall charge solution should be made more global by looking for redundant blobs over projections involving the time dimension (equivalently, the drift dimension). In this manner, deghosting and charge solving can be seen as complementary steps, each aiming to remove redundant solutions (ghosts) across different marginalizations of the full 3D charge distribution.

When viewed in 3D, adjacent blobs across time slices can be connected to form proto-clusters. These form an intermediate step towards a complete description of the activity in an interaction; however, gaps between proto-clusters within an interaction may exist because of both real and artificial causes, such as π^0 decay and poor reconstruction respectively. The step of forming proto-clusters from blobs can be thought of as a form of tiling over the time (drift) dimension. Some of these proto-clusters are ghosts, and are redundant in explaining the charge measured across wires. To identify these ghosts, their structure and relation to non-functional regions and real charge are examined.

First, as is typically the case, ghosts will often exist in detector regions with a non-

functional wire plane. Ghosts must lie along at least two sets of fired wires, corresponding to coincidental charge distributions at other locations. Track-like ghosts can hide in the non-functional region of one wire plane, and in the projected shadow of real charge distributions along the other wire planes. However, as the 2D anode-projected position of charge has already been solved under the constraint of sparsity through the imaging equation, it is likely that ghosts hide, and therefore are redundant, in all three 2D wire-versus-time projections.

This motivates the deghosting procedure, where each proto-cluster is examined for redundancy by considering whether the charge distribution minus the candidate proto-cluster, viewed under each 2D wire-versus-time projection, is capable of describing the set of fired wires. Because the removal of ghost proto-clusters decreases the total charge of the 3D distribution, the imaging equation must be re-solved to recover an accurate charge distribution. The remaining blobs after deghosting are used as input for this round of charge solving. This overall procedure is iterative, and empirically it was found that three rounds of charge solving plus deghosting were useful, with further rounds providing minimal benefit.

5.2.4 Clustering

The proto-clusters established in the previous section help to group the activity of a particle within an interaction, but they are disjointed and insufficient for future reconstruction goals. It is necessary to group all of the activity of an event together, for organizational purposes and algorithmic ones. In particular, charge-light matching in Sec. 5.3 requires a full count of all the charge in an event to be able to accurately match this charge with the corresponding light signal it produces. Although charge may reach the TPC at different times within the drift window, the true interaction occurs at relativistic speed and so all parts of an interaction begin emitting scintillation light to be recorded by the PMT system simultaneously. Therefore, light activity is recorded as a single flash, and charge-light matching will be most successful if the entire interaction can be grouped into a single charge cluster.

Existing proto-clusters are incomplete for a number of reasons. There are gaps resulting from non-functional regions (3% of the detector volume) as well as errors in signal processing, particularly in the case of difficult topologies such as isochronous and prolonged tracks. A cluster may also contain multiple separated components resulting from non-ionizing neutral particles that are invisible to the LArTPC but may later decay to ionizing particles, such as $\pi^0 \rightarrow \gamma\gamma$ and $\gamma \rightarrow e^+e^-$.

It is also possible that an interaction may be over-clustered, where charge from separate coincidental interactions overlap. The chance of this happening is largely driven by the existence of cosmic rays, which populate every drift window, although usually without overlapping. The source interactions for overlapping charge occur at different times, and therefore produce separate light flashes, but it is possible for their ionized charge to reach the wire planes simultaneously, hence the overlap. This can happen, for example, if one interaction occurs early in a drift window near the cathode while a second interaction occurs later in time but near the anode. It is important to de-cluster these interactions both to accurately describe each interaction for purposes of particle identification and event selection, and to allow for accurate charge-light matching. Finally, all of the sources of imperfection discussed can hamper the effectiveness of charge solving and deghosting, producing some ghosts that cannot be removed through the methods discussed earlier, so some residual ghosts remain as isolated proto-clusters.

Recovering the original structure in the presence of gaps is based on two metrics: the distance between the gaps and the directionality of both proto-clusters. The distance vector between clusters is defined as the minimum distance between pairs of points across the clusters, and is computed by constructing a k-dimensional graph connecting the 3D charge points through nearest neighbors using the nanoflann package [112]. The direction of each cluster is determined by taking the Hough transform [113] of the point cloud, which maps the 3D position of each point to multiple locations on a 4D phase space characterizing the

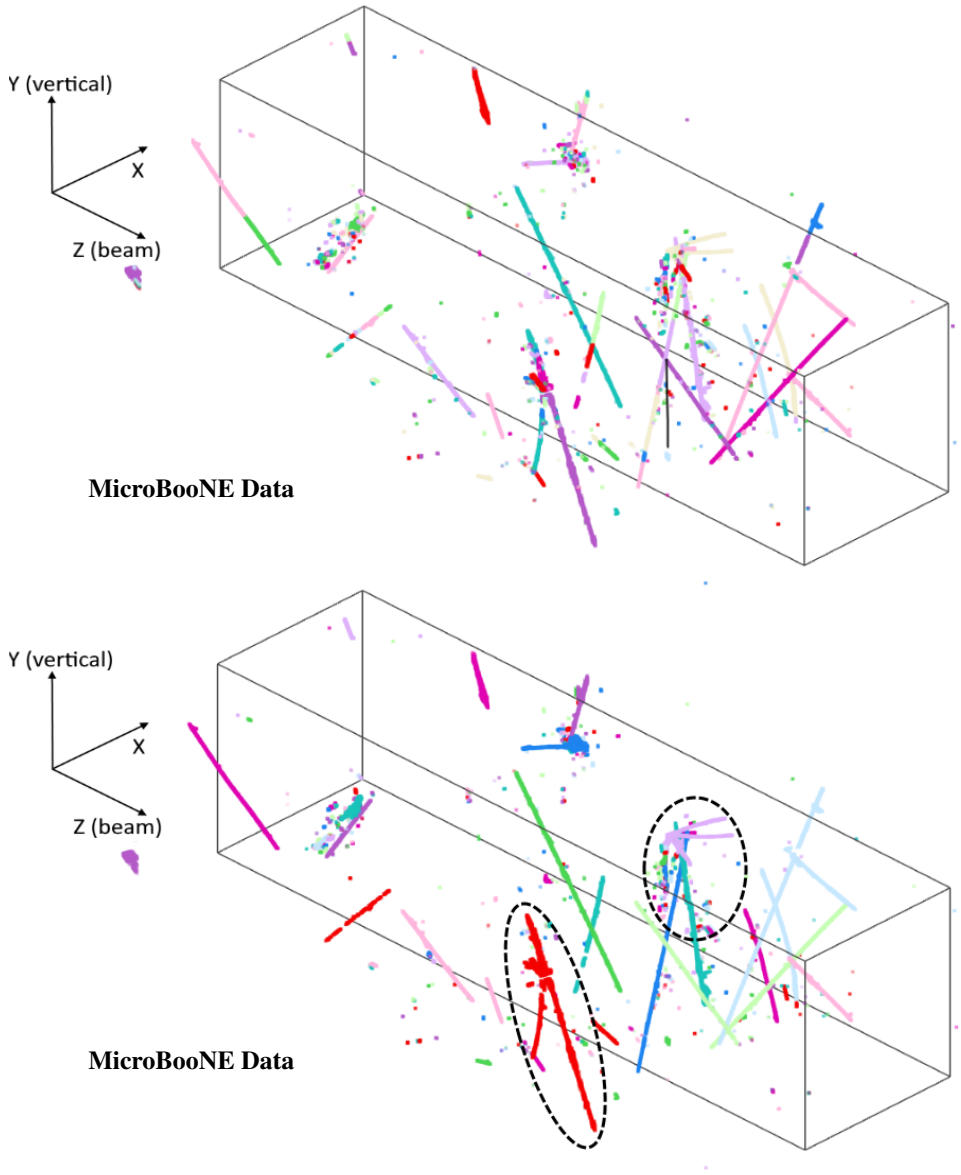

Figure 18: Demonstration of the effectiveness at bridging gaps in clusters. The solid black box represents the LArTPC active volume with an x-position (converted from the readout time) relative to the neutrino interaction time. Top: proto-clusters solely based on proximity. Bottom: clusters after bridging gaps. The two circles indicate remaining clustering issues, e.g. over-clustering of cosmic-ray muons and under-clustering of neutrino interactions. Cluster membership is indicated by uniform color. Source: [106]

position and orientation of all possible lines. If the space of possible line orientations, $SO(3)$, is parameterized by n possible orientations, then each 3D space point will map to n locations in 4D phase space, where two dimensions controlling the physical offset of the line are fixed

by the 3D location of the point in the domain. This can be thought of conceptually as taking the plane normal to the orientation being considered and projecting all 3D points onto this plane and recording their 2D projection. The set of n orientations under consideration are then iterated to complete the mapping. For a given orientation, points along nearby parallel lines will map to nearby phase space locations. If a point cloud is track-like, only the orientations that match the directionality of the point cloud will map all 3D points into a dense region in phase space. The best orientation can be easily found by searching for regions of highest density in the phase space. Clusters are merged based on the distance between them as well as the difference in directionality between them. The performance of bridging gaps in clustering is demonstrated in Fig. 18.

After clusters are merged across gaps, over-clustering from coincidental overlaps is detected and the clusters are separated. Candidates for over-clustering are found by inspecting the Principle Component Analysis (PCA) [114] of the point cloud of charge. Track-like events will only have a single large PCA eigenvalue, while overlapping tracks will have two non-trivial PCA eigenvalues. Over-clustered events are detected by looking for a primary track that extends along the length of the largest PCA axis. This is achieved by constructing the smallest convex shape bounding the point cloud using the quickhull algorithm [115], and looking for endpoints indicative of the primary track near the boundary of the convex region. These endpoints must be located near the PCA primary axis, and the local point cloud must extend along this axis. Once a primary track of an over-clustered event is selected, a Kalman filter [116] is used to crawl along the track to locate the other end. Then, Djikstra's shortest path algorithm [117] is run over the graph of charge points to assign a simple trajectory along the track, and the local charge points along the trajectory are associated with it. These points are considered one cluster, and are removed from the point cloud. This process can be iterated to further separate clusters in the case of multiple overlaps. The performance of separating over-clustered events is demonstrated in Fig. 19.

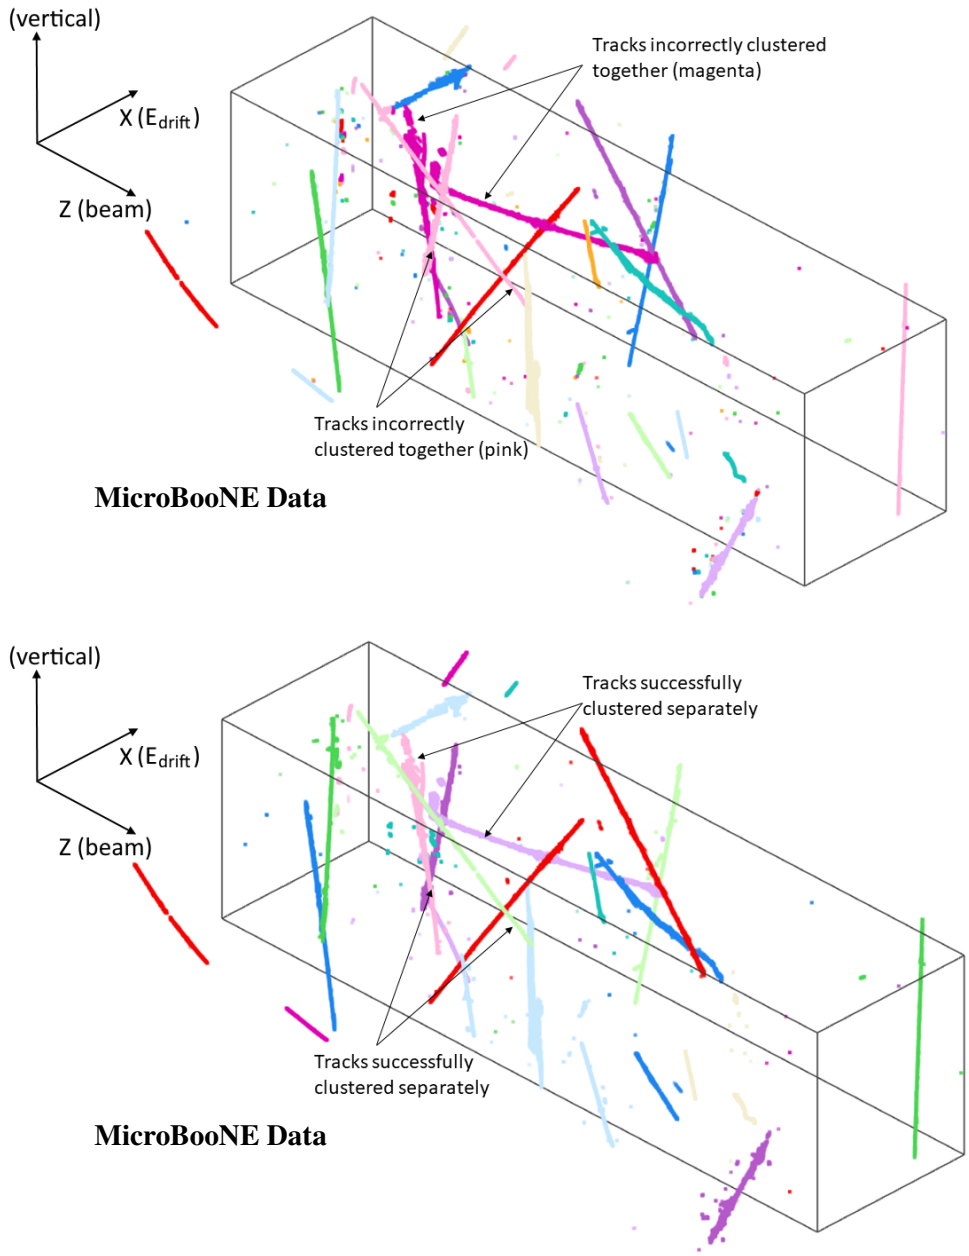

Figure 19: Demonstration of the effectiveness of the clustering algorithm to separate coincidental overlaps in clusters. The solid black box represents the LArTPC active volume with an x-position (converted from the readout time) relative to the neutrino interaction time. The top and bottom panels show the clusters before and after applying this algorithm. Cluster membership is indicated by uniform color. Source: [106]

After clustering to bridge gaps and separate coincidental overlaps, there is room for improvement with another round of deghosting. By connecting proto-clusters across gaps,

ghosts that could not previously be found redundant within a single proto-cluster can now be identified when the full cluster information is available. After coincidental overlap separation, some ghost proto-clusters are now properly isolated and may be removed without also removing sections of valid charge.

There is a final clustering algorithm dedicated to group isolated components in a neutrino interaction, which can have more complicated topologies, such as $\pi^0 \rightarrow \gamma\gamma$ decay, than cosmic ray activity. This algorithm is only run on the 2.3 ms of drift window following the beam spill, corresponding to candidate neutrino interactions. It follows a similar technique to the gap-bridging algorithm, where the directionality of a (> 1 cm long) cluster is determined through use of the Hough transform. The trajectory of the cluster is extended and examined to see whether it intersects with the endpoint or extended trajectory of another cluster. These cases indicate a common decay vertex, although not necessarily the neutrino vertex, and all associated clusters are merged.

The clustering achieved in this section, building on the broader suite of imaging algorithms before it, finishes the task of preparing a robust 3D description of the data for each TPC interaction. The fully-3D charge point cloud solved for in Sec. 5.2.2 represents a novel step forward in LArTPC physics and enables more powerful implementations of high-level reconstruction algorithms by virtue of working with the more physical 3D charge distribution instead of 2D projections. Deghosting in Sec. 5.2.3 fixes many of the problems introduced by taking on the ambitious goal of 3D imaging under a 2-plane requirement, producing a solved charge distribution with minimal errors. Finally, this distribution is assembled into completed clusters corresponding to entire interactions in Sec. 5.2.4 so that the results produced may be easily fed into upcoming algorithms.

5.3 Charge-Light Matching

The completion of charge reconstruction through imaging marks the transition from processing data into fundamental, information rich descriptions toward the creation of high-level descriptions of the particles that produced these interactions. In-between these two halves of the reconstruction chain sits flash-light matching. The clusters produced in Sec. 5.2.4 enable the pairing of TPC charge activity with the corresponding light flash as measured by the PMT array. This has the primary benefit of allowing the precise interaction time to be linked with the charge cluster, as the PMT array is capable of measuring prompt scintillation light at sub-μs resolution. Precise interaction timing allows cosmic ray background events to be rejected using the narrow beam window (1.6 μs) instead of the lengthy drift window (2.3 ms). There are additional benefits derived from knowing the drift offset: cosmic ray data can be studied to form the effective detector boundary described in Sec. 4.3.4, and candidate neutrino position will be used in Sec. 6 to remove cosmic rays.

5.3.1 Flash Reconstruction

The light collection system consists of 32 PMTs arranged across the anode to maximize coverage of the TPC volume. They measure the light produced over a $\sim$ 23.4 μs window, beginning 4 μs before the beam spill and consisting of 1,500 samples taken at a rate of 64 MHz. There is also a cosmic discriminator window that triggers on light activity throughout the entire 6.8 ms PMT readout window, and stores data for 40 samples, corresponding to 0.6 μs. A Fast Fourier Transform (FFT) [118] algorithm is used to deconvolve the raw waveform to unfold the electronics response. Flashes are identified based on the paired requirements of *a*) at least three PMTs measuring $\geq$ 1.5 PE and *b*) more than 6 PE measured in total, both evaluated based on the total PE measured over a rolling 100 ns (7 samples) window. Once a flash is detected, the flash start time is set as the sampling bin with the largest total PE, and any PE activity measured in the following 7.3 μs is associated with the flash, to include the

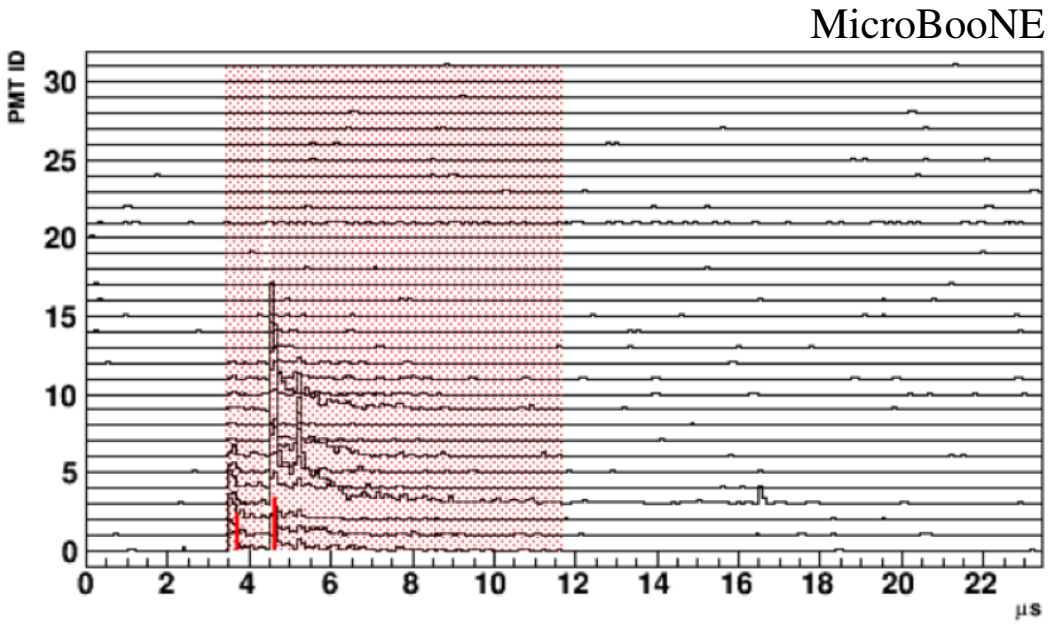

Figure 20: Deconvolved waveform of the 32 PMT array. The red bands represent flash windows for two separate flashes, at times of 4.6 μs and 5.3 μs. The measured PE is represented by the z-axis. Source: [106]

late scintillation light contribution. However, if another flash is detected before the end of the first flash window, as is the case in Fig. 20, it is cut short and the remaining PE is associated with the later flash. To trigger this, there must be a spike in the PE profile over time, and it must either be more than 1.6 μs after the original flash time or it must produce a significantly different PE distribution across the PMTs as determined by a Kolmogorov-Smirnov (KS) [119] test.

5.3.2 Match Determination

In the entire PMT readout window there are typically 30 – 50 light flashes reconstructed, as shown in Fig. 21. This is larger than the ~ 30 TPC charge clusters recorded because the PMT array is capable of detecting light from activity outside the fiducial volume but still inside the cryostat. As a result, some flashes have no corresponding charge cluster. However, there are also charge clusters with no corresponding reconstructed flash, particularly in cases of

low visible energy interactions and those occurring near the cathode (away from the PMTs). Additionally, there are some inaccuracies in the clustering in Sec. 5.2.4, resulting in cases where multiple clusters should be grouped together, and so correspond in reality to a single flash. This motivates a charge-light matching algorithm that is capable of handling the variety of possibilities instead of exclusively or exhaustively pairing flashes to clusters.

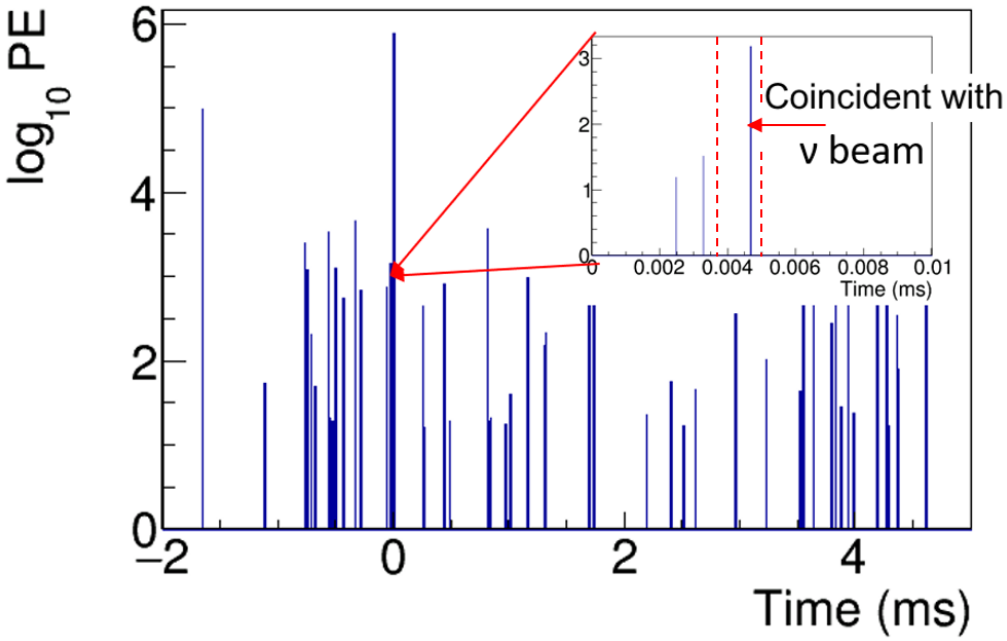

Figure 21: Total reconstructed PE as a function of time across the entire 6.4 μs PMT readout window. In general, there are 40–50 reconstructed PMT flashes in each BNB event. Source: [106]

Matches are determined by comparing the measured flash on the PMT arrays against a predicted light distribution based on the measured charge in a cluster. This predicted flash is computed using a photon library [120] mapping charge across 3D voxel locations to predicted light at the PMT array based on a large number of simulations. To accurately predict a flash measurement, under each candidate charge-light pairing the charge cluster is moved along the drift (x) direction to the location set by the flash time, adjusting the light prediction at the location of the PMTs. Before proceeding to a global fit to match flashes to charge clusters, a pre-selection is made to reduce the number of matches considered, and therefore reduce the

complexity of the fit. The pre-selection restricts the set of clusters that can be matched to a flash based on two requirements. First is a requirement of plausible causality: at the time of the flash, the charge cluster must be fully contained within the TPC volume, accounting for drift position. Second, the predicted and measured light distributions must be roughly compatible, as determined by both a KS test and a comparison of overall PE.

Now that many of the potential matches are eliminated, the global fit is performed. The match relations are formulated in a similar manner to the imaging equation, relating PMT-measured light to charge-based predictions via $y = Ax$. Again, this is an under-determined system with more unknowns than knowns, and the method of compressed sensing is used to arrive at a probable solution through the use of an ℓ_1 norm regularization term R_1. There are also regularization terms R_2 and R_3 enforcing the constraint that each cluster have only one match and that the number of unmatched flashes should be small, respectively. The total χ^2 test statistic is given as:

$$\begin{aligned}
\chi^2 &= \sum_i \sum_j \chi^2_{ij} + R_1 + R_2 + R_3 \\
\chi^2_{ij} &= \frac{\left(M_{ij}(1-b_i) - \sum_k a_{ik} P_{kij}\right)^2}{\delta M^2_{ij}} \\
R_1 &= \lambda_1 \sum_i \sum_k a_{ik} \\
R_2 &= \lambda_2 \sum_i \left(\sum_k a_{ik} - 1\right)^2 \\
R_3 &= \lambda_3 \sum_i b_i
\end{aligned} \tag{5.14}$$

where the indices i, j, and k run over flashes, hit PMTs in a flash, and charge clusters, respectively. M_{ij} and δM_{ij} represent the measured charge and its uncertainty on the j'th PMT of the i'th flash, respectively. P_{kij} represents the predicted light from the k'th cluster at the i'th flash time on the j'th PMT. $a_{ik}, b_i \in [0,1]$ represent the degree of certainty that flash i matches with cluster k and a dummy parameter that allows for a flash to find no

match, respectively, and are solved for in the χ^2 minimization. Finally, λ_1, λ_2, and λ_3 are regularization strengths.

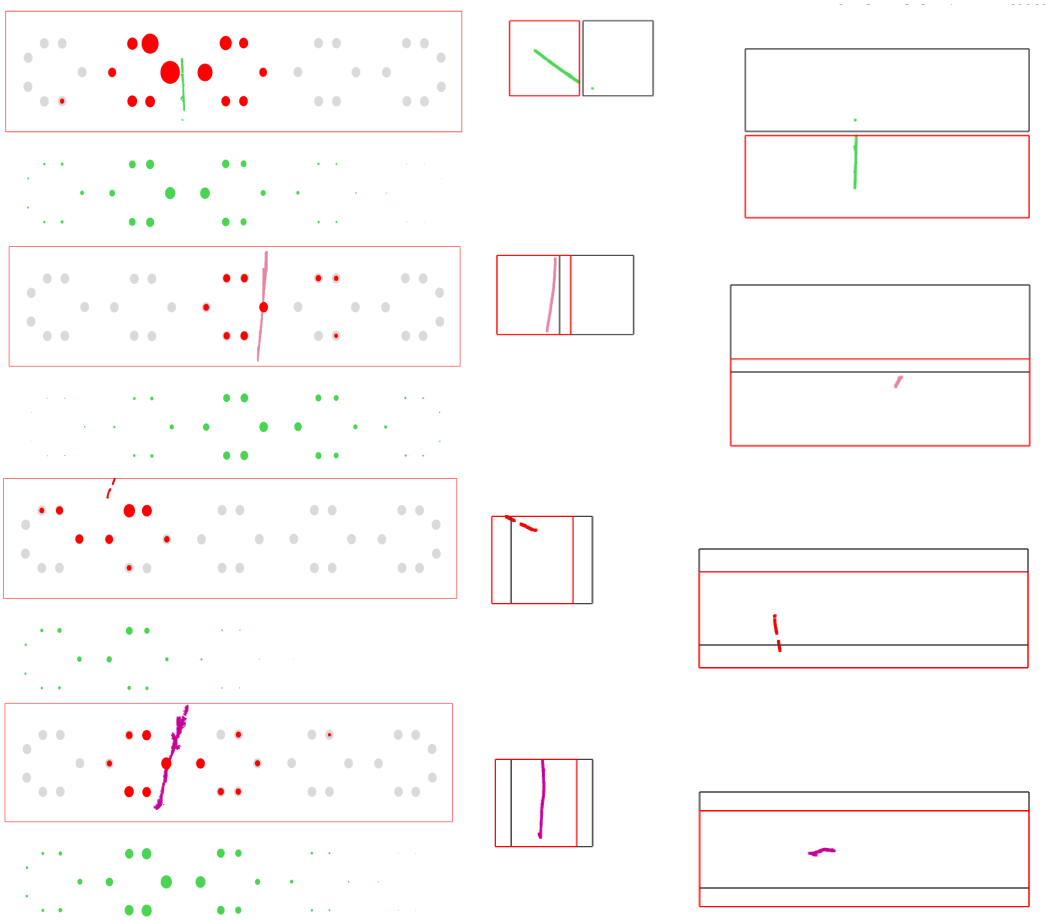

Figure 22: A selection of 4 out of 31 matched pairs, corresponding to cosmic rays, from a data event readout window. The y-z (left), y-x (center), and x-z (right) 2D projections of the charge distribution in the detector are shown, with the matched drift window in red and the neutrino drift window in black. In the y-z view, grey circles show the PMT locations on the anode, with red circles representing the magnitude of measured PE. Offset below for visibility, the predicted light of the matched cluster is shown. Source: [106]

After the fit is performed, a number of flashes will remain unmatched, with $a_{ik} \approx 0$. These are removed from the set of considered flashes, all b_i values are fixed at 0, and a second round of minimization is performed to solve for a_{ik} alone. Afterwords, each cluster is examined, and the strongest flash match as determined by a_{ik} is selected. Since each flash is allowed to match to multiple clusters, the flashes and their corresponding matched

clusters are examined. For each flash, the largest matched cluster is considered the principle component, and subsequent cluster matches are considered, and kept if they improve the match compatibility, as defined by the KS test and overall PE comparison described earlier. The unmatched clusters are tested against unmatched flashes to look for any remaining matches that have been missed.

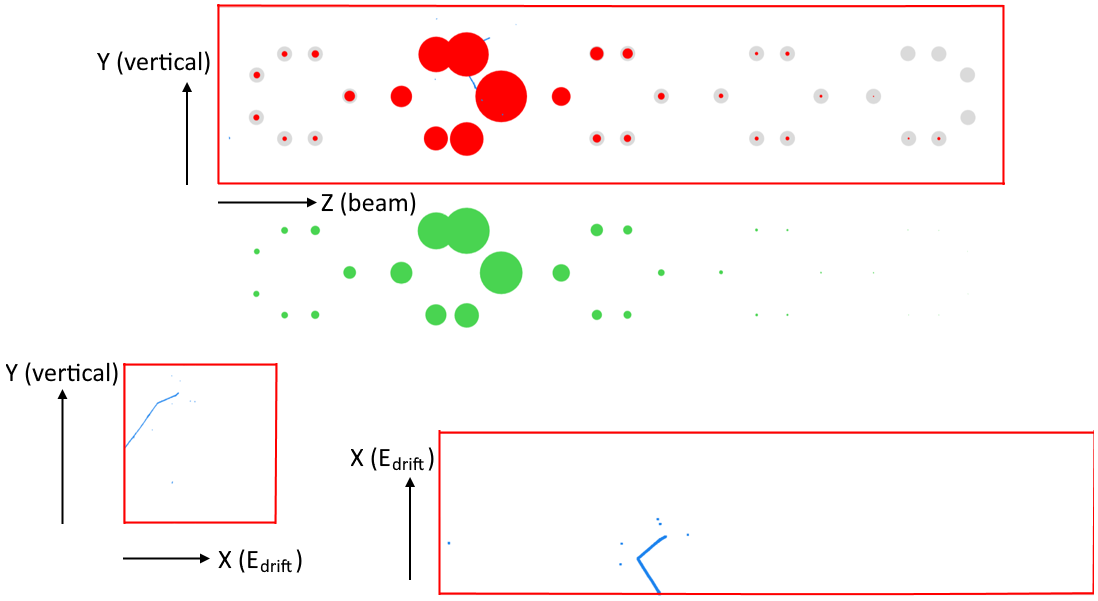

Figure 23: A matched muon neutrino interaction from a data event. The y-z (left), y-x (center), and x-z (right) 2D projections of the charge distribution in the detector are shown, with the matched drift window in red. In the y-z view, grey circles show the PMT locations on the anode, with red circles representing the magnitude of measured PE. Offset below for visibility, the predicted light of the matched cluster is shown. Source: [106]

Fig. 22 shows 4 out of 31 matched flash-cluster pairs, corresponding to cosmic rays, from a data event readout window, including the comparison between measured PE in red and predicted PE in green. Fig. 23 shows the match for a muon neutrino interaction. Flash matching is the first step in a series of event selection algorithms, as will be discussed in Sec. 6, and is also an important part of the reconstruction chain. Through matching, the results of clustering are improved, and the drift distance is established. Matching is computationally important as it represents a $\sim 40\times$ reduction in candidate neutrino events, owing to the removal of non-beam-coincident cosmic rays, significantly reducing the amount

of work generated in upcoming reconstruction algorithms. Finally, he use of a many-to-many matching strategy achieves a significant performance improvement over one-to-one solutions that aim to only match candidate neutrino flashes, leaving the majority of cosmic rays unmatched. Not only does matching cosmic rays assist other areas of the analysis, such as effective detector boundary mapping, but it also improves the neutrino flash match accuracy. Through the use of a global fit, information from cosmic ray flashes is incorporated to help reduce the chance of accidentally matching neutrino flashes to cosmic ray charge clusters, preserving a higher selection efficiency. The performance of charge-light matching in producing a high-quality neutrino selection is discussed in more detail in Sec. 6.

5.4 Trajectory Fitting

Trajectory fitting is a multi-purpose procedure. The primary and immediate goal is to aid the identification of cosmic rays through a computation of the charge deposition rate dQ/dx of a particle track. However, the benefits are far larger. Through dQ/dx, the directionality of many particle tracks can be determined. Also, measurements of dQ/dx translate easily to energy deposition dE/dx, which is useful for particle identification and energy estimation. More generally, in computing the trajectories of each particle in an interaction a graph structure of the entire interaction (or at least its connected components) is constructed, connecting the flow between various particles. This makes the future steps easier, as half the work is already done in determining which particles decayed to which. The larger construction of particle flow as a top-level description of an entire neutrino interaction is a logical extension of the groundwork built in trajectory fitting. The trajectory fitting algorithm discussed in this section is focused on determining the trajectory of a single particle; in Sec. 5.5 it will be expanded upon to accurately parse each individual track in a neutrino interaction.

The task of global trajectory fitting is largely enabled by a 3D reconstruction of the charge distribution, as determining a seed trajectory would be far more challenging without the

physically relevant 3D point cloud. However, there are still numerous challenges that remain, and drive the overall shape of the procedure. There are inaccuracies in the 3D charge points, especially resulting from smearing through the Wiener filter in signal processing as well as the challenging reconstruction of difficult topologies (isochronous, prolonged, and aligned tracks). As always, regions with non-functional wires reduce the reconstruction quality in their vicinity. There is also an issue particular to trajectory fitting, the difficulty in fitting a trajectory because of the non-linear nature of the fit. Simply feeding the global point cloud distribution into a test statistic minimization calculation would be highly unstable and fail to produce a reliable trajectory [95].

For these reasons, the process of determining a trajectory is split into multiple sections. First, a seed trajectory is established through a graph-based treatment of the point cloud of charge. The seed trajectory is intended to be reliably accurate in describing a particle through the general path it takes, although the exact positioning of its points may be wrong. From this, a precision trajectory in 3D is computed by projecting the 3D trajectory points into each 2D (wire number vs time) plane projection and simultaneously fitting to the pixels within each projection, inspired by the Projection Matching Algorithm in ICARUS [121]. The use of 2D projections in the fit avoids the issues with 3D points result from track topology and signal processing and are limited to one plane projection. Meanwhile 3D trajectory points still provide crucial information on the geometry of the interaction and on relating these 2D projections. From trajectory fitting an ordered set of 3D points $S\{x_i, y_i, z_i\}$ with associated 2D projections $U\{u_i, t_i\}$, $V\{v_i, t_i\}$, $W\{w_i, t_i\}$ is created. Once trajectory points are established, the measured charge is assigned along the track to fit a dQ/dx profile. By separating this step from the prior trajectory fit, issues resulting from the non-linear nature of the overall fit are minimized to allow for a stable procedure.

5.4.1 Minimum Spanning Tree

The trajectory seed is established in a robust manner by first organizing the 3D charge point cloud into a graph capable of accurately describing the relation of points along a particle trajectory. In theory a simple graph based on nearest-neighbor connections could be formed as was used in Sec. 5.2.4. This would be insufficient for the level of precision desired here, because such a graph could essentially be over-connected. Intuitively, the problem can be seen when assigning a trajectory around a hard turn such as at a vertex. In a highly connected graph, the shortest path from one end to another would cut corners, failing to well describe particle trajectories. This is especially important under the very same example of particle vertices, where vertex placement can play a critical role in accurate particle identification.

To prevent a shortest-path algorithm from deviating from real trajectories in such a manner, the graph used to connect point cloud points is based on a Steiner tree [122]. A Steiner tree takes as input an unordered graph of vertices and edges with non-negative weights, where a subset of the vertices are labeled as Steiner terminals. The tree is then constructed as a minimum spanning tree connecting these terminals, so that the sum of the weights used in a minimum. It may include connections to other vertices, but has no requirement to connect them at all, making it a sub-graph of the original. Steiner terminals are selected as local maxima charge points (higher charge than any nearest neighbors) above a threshold charge. By selecting high-charge points as Steiner terminals, the shortest path is required to follow along these points. Naturally, high charge points tend to be concentrated around the center of a trajectory, and so the shortest path over the Steiner tree much more closely follows actual particle trajectories. Edges of the graph are computed by scaling the Euclidean distance by a function based on their charges to preference connections between high-charge points. Since forming the Steiner tree is NP-complete and computationally slow, a Steiner tree greedy algorithm [123] is used to approximately solve it by connecting nearby regions as determined through a Voronoi tiling. Results are shown in Fig. 24, comparing the

performance before and after weighting edges based on charge magnitude and uncertainty.

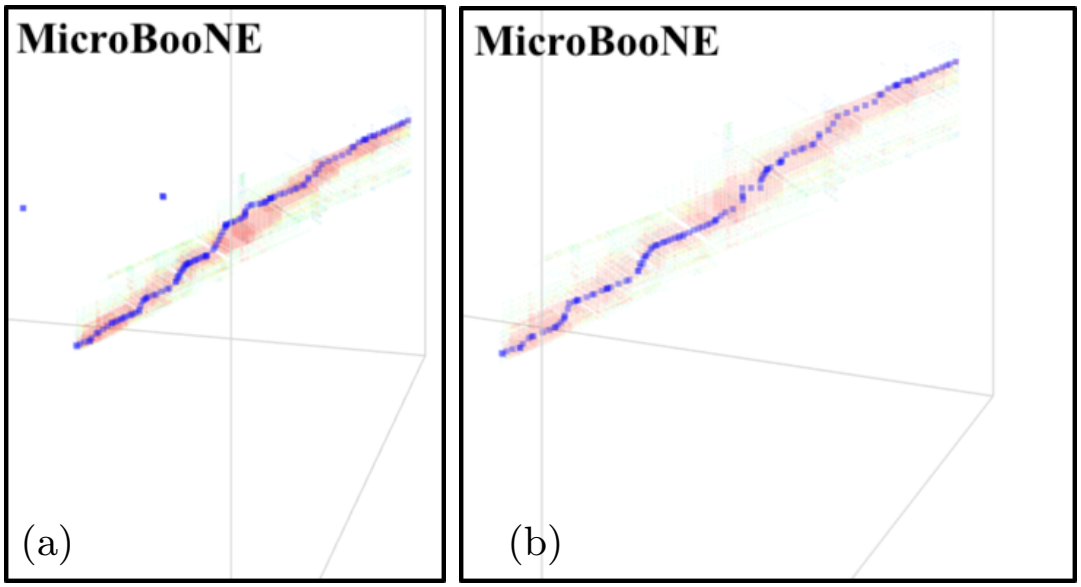

Figure 24: Performance of Steiner-tree-based shortest path algorithm before (a) and after (b) including charge weights to otherwise Euclidean graph edges. The shortest path forms the basis of the trajectory seed to be fit. Source: [95]

5.4.2 Trajectory Seed

The trajectory seed is computed by connecting the extreme points on either end of the point cloud (as determined using PCA) through the shortest path along the Steiner tree. The ordered set of graph points along this path form the basis of the trajectory seed. To improve the upcoming trajectory fit, points are added or removed as needed along the path to maintain a spacing of ~ 1 cm between points. In preparation for the trajectory fit, the seed points are associated with the nearby 2D (wire number vs time) pixels. This association is formed for each trajectory seed point by considering the nearby 3D points, finding their parent blobs, and projecting into the 2D views to locate close-by 2D pixels. By fixing the trajectory-point to pixel associations, the computation of the fit can be controlled to avoid fit

instability.

5.4.3 Trajectory Fit

To compute the trajectory fit, a test statistic T is constructed, comparing the fit quality of each 3D point projected into each 2D view:

$$T(S\{x_j, y_j, z_j\}) = \sum_{k=u,v,w} T_k \tag{5.15}$$

where k is an index over the U, V, and W wire planes and $S\{x_j, y_j, z_j\}$ is the set of trajectory points to be fit. The test statistic for each plane is computed as:

$$T_k = \sum_j \sum_i \frac{q_i^2}{\delta q_i^2} \cdot (\Delta L_k)_{ij}^2 \tag{5.16}$$

$$(\Delta L_u)_{ij}^2 = \Delta u^2 \cdot (u_i - u_j(y_j, z_j))^2 + \Delta x^2 \cdot (t_i - t_j(x_j))^2 \tag{5.17}$$

where the construction of ΔL_v and ΔL_w follow the example shown for ΔL_u, and each represent the distance between the projection of the j'th 3D trajectory point and the i'th pixel. q and δq represent the measured charge and its uncertainty at the associated pixel on a wire plane readout, enhancing the weight of high-charge points and suppressing the weight of high-uncertainty points. Δx and Δu represent the width of the time slice (2 μs, corresponding to 2.2 mm of drift) and wire pitch (3 mm) respectively. The indices i and j run over nearby 2D pixels and 3D trajectory points respectively. The 2D coordinates (wire number, time slice) for pixel i are given by u_i and t_i, respectively, and correspondingly u_j and t_j are the projected coordinates from the j'th 3D trajectory point. The projection from 3D $\{x_j, y_j, z_j\}$ coordinates

to $\{t_j, u_j, v_j, w_j\}$ is computed as:

$$\begin{aligned} t &= \frac{1}{\Delta x} \cdot x + t_0 \\ u &= \frac{1}{\Delta u} \cdot (-sin(\theta_u) + cos(\theta_u) z) + u_0 \\ v &= \frac{1}{\Delta v} \cdot (-sin(\theta_v) + cos(\theta_v) z) + v_0 \\ w &= \frac{1}{\Delta w} \cdot (-sin(\theta_w) + cos(\theta_w) z) + w_0 \end{aligned} \tag{5.18}$$

where θ_u, θ_v, and θ_w are the wire orientations with respect to the vertical direction for each wire plane, and $\{t_0, u_0, v_0, w_0\}$ are the coordinates of the origin. These equations can be combined to form an overall matrix equation for the test statistic:

$$T = \sum_{k=u,v,w} (M_k - P_k \cdot S)^2 \tag{5.19}$$

where M_k is the charge-weighted 2D pixel coordinates and P_k is the charge-weighted projection matrix for each wire plane. Minimizing T leads to a matrix equation to solve for S:

$$\left(\sum_{k=u,v,w} P_k^T P_k \right) \cdot S = \left(\sum_{k=u,v,w} P_k^T M_k \right) \tag{5.20}$$

The trajectory S is solved for numerically by using the Biconjugate Gradient Stabalized method (BiCGSTAB) [124].

5.4.4 dQ/dx Fit

Since the trajectory has already been determined, the hard work is done and dQ/dx can easily be computed by minimizing a test statistic T' that compares the projected charge from each trajectory point to the charge of associated 2D pixels:

$$T'(S\{Q_j\}; S\{x_j, y_j, z_j\}) = \sum_{k=u,v,w} T'_k + T'_{reg} \tag{5.21}$$

$S\{Q_j\}$ is an ordered set of charges for each corresponding trajectory point. The regularization term promotes a smooth solution by computing a penalty based on the second derivative of the dQ/dx distribution over trajectory points:

$$T'_{reg} = \sum_i \left(\sum_j F_{ij} \cdot \frac{Q_j}{s_j} \right)^2 \tag{5.22}$$

where s_j represents the length of the j'th segment, computed as the average distance to each adjoining point. This makes $\frac{Q_j}{s_j}$ an approximation of dQ/dx at j. F encodes the second derivative information by taking the difference between adjacent $\frac{Q_j}{s_j}$ entries when the product is summed:

$$F = \eta \cdot \begin{bmatrix} -1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & -2 & 1 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 1 & -2 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix} \tag{5.23}$$

where η provides an overall normalization strength, set to 0.3 or 0.9 when an adjacent non-functional wire is on the induction or collection plane respectively. The use of regularization is important to dampen the impact of poorly defined trajectory points in the vicinity of non-functional wires. The test statistic over each plane is computed as:

$$T'_k = \sum_i \frac{1}{\delta q_i^2} \cdot \left(q_i - \sum_j R_{ij}^k Q_j \right)^2 \tag{5.24}$$

where R_{ij}^k computes the detector response from original ionization charge along the trajectory Q_j and the measured charge q_i at a wire. Because of computation time costs, instead of simulating the full chain of detector effects, including diffusion, TPC field response, electronics response, and signal processing, a much faster approach is used to approximate an effective signal response. A Gaussian distribution is used, with width computed in quadrature

from contributions of diffusion and the software filter:

$$\begin{aligned}
\sigma_{D_L} &= \sqrt{2D_L \cdot t_{drift}} \\
\sigma_{D_T} &= \sqrt{2D_T \cdot t_{drift}} \\
\sigma_{F_t} &= 1.57\,\text{mm} \\
\sigma_{F_u} &= 0.36\,\text{mm} \\
\sigma_{F_v} &= 0.66\,\text{mm} \\
\sigma_{F_w} &= 0.11\,\text{mm}
\end{aligned} \tag{5.25}$$

where D_L and D_T are computed in Eqn. 4.7, and t_{drift} is the drift time.

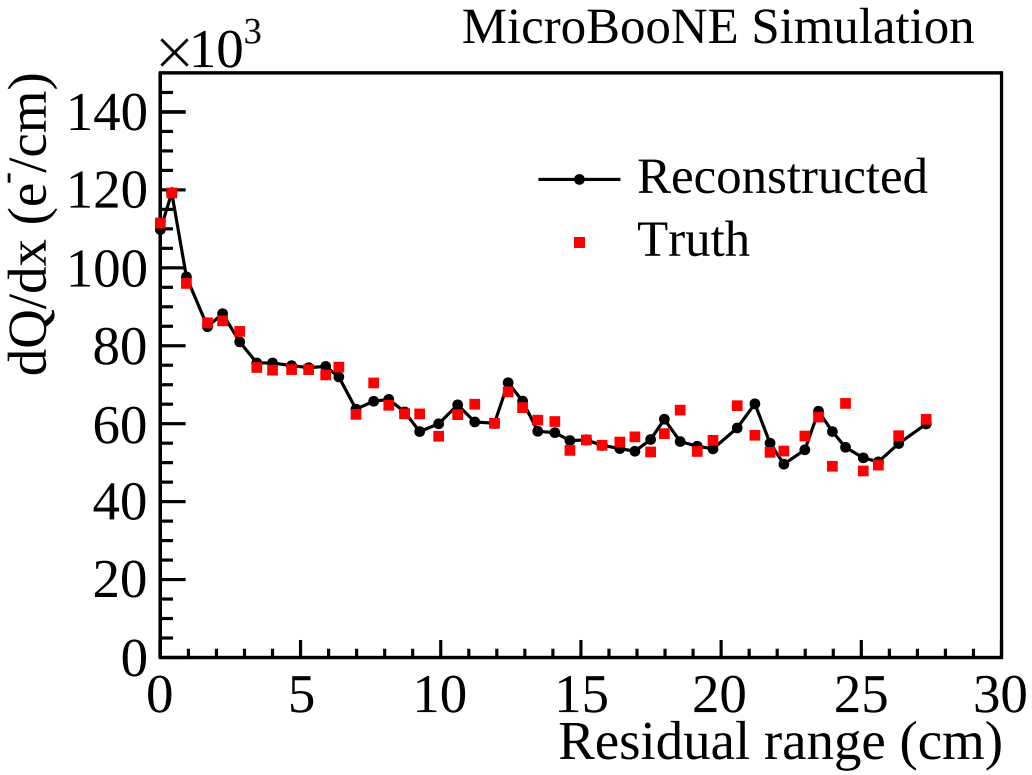

Figure 25: The best-fit dQ/dx (in black) along a simulated muon track trajectory is compared with the true dQ/dx (in red) as a function of the residual range (distance along the track with respect to the stopping location). Source: [95]

The fit dQ/dx as a function of residual range (distance along the track with respect to

the stopping location) demonstrates the capability to detect particle directionality, as shown on simulation in Fig. 25. From the Bethe-Bloch equation 5, the signature Bragg peak can be seen as the muon comes to rest. Determining particle directionality is an important result in general, for determining vertices of origin and overall particle flow. It is also particularly useful for identifying cosmic ray muons to efficiently remove them from a neutrino selection. Cosmic rays typically enter from above and outside the detector, providing a signature that can be identified from determining a particle's trajectory.

5.5 Particle Identification and Vertexing

This section completes the Wire-Cell reconstruction chain by separating a neutrino interaction into a tree that tracks the flow of individual particles. Each particle has its trajectory computed in a manner built upon the algorithm described in Sec. 5.4, and vertices are created to label the decay/origin location of particles. Finally, the neutrino vertex is located and the overall neutrino energy is computed [125]. Note that these algorithms are only run on candidate neutrino interactions that pass a number of event selection cuts based on the flash-light matching and simple particle trajectory algorithms of Sec. 5.3 and Sec. 5.4.

5.5.1 Track Segmentation

Track segmentation expands on the previously established single-trajectory fit algorithm, which only follows the primary (longest) track in an interaction for the purposes of cosmic ray muon identification. Taking the endpoints of this trajectory as an input, a Kalman filter [116] is used to crawl along the trajectory and look for kinks, identified by sharp turns in the trajectory. Wherever a kink is located, the track is split into two segments and a vertex is created at the intersection. Each segment is re-fit using the trajectory fit algorithm to determine a new trajectory and dQ/dx. Those fit trajectories are removed and the remaining cluster is iterated over to look for additional segments to be fit. In this manner, all track

segments in the interaction are identified and fit. At the end, and vertices in close proximity to each other are merged. However, the determination of precise vertex location is left for the next section. Segmentation and track fitting performance is demonstrated in Fig. 26.

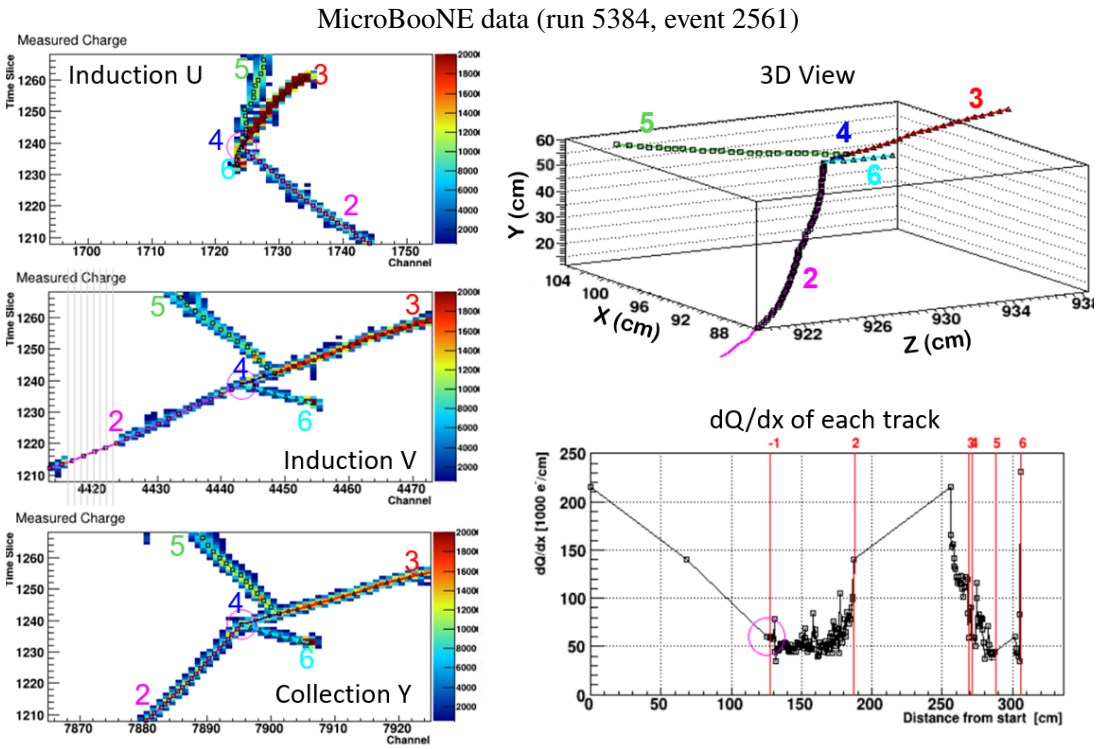

Figure 26: Example of multi-track trajectory fitting. The 2D projections on the U, V, and W planes are shown on the left, and a 3D view is shown in the top right. The fit dQ/dx of each trajectory is shown in the bottom right. Source: [125]

5.5.2 Vertex Fitting

Vertex fitting is a fully 3D algorithm, relying on the trajectory points for each track segment connected to a vertex. For each segment, iterated with the variable i, the trajectory points near the vertex are considered, omitting those within 1.5 cm as they may be associated with the wrong track segment. PCA [114] is run over the selected trajectory points to establish the eigenvectors $\{\vec{v}_{i1}, \vec{v}_{i2}, \vec{v}_{i3}\}$ and eigenvalues $\{\lambda_{i1}, \lambda_{i2}, \lambda_{i3}\}$ that describe the variance in the trajectory, with $\vec{v}_{i1}$ describing the overall track trajectory.

The distance from the center of each set of trajectory points to the vertex is described by $\vec{r}_i$. Then, the minimum distance from the vertex to the primary axis of a track (along $\vec{v}_{i1}$) is described by:

$$s_i^2 = (\vec{r}_i \cdot \vec{v}_{i2})^2 + (\vec{r}_i \cdot \vec{v}_{i3})^2 \tag{5.26}$$

The vertex location is computed by minimizing a test statistic defined based on a weighted computation of s_i^2, as well as a sort of regularization term:

$$T = \sum_i \lambda_{i1} \left(\frac{(\vec{r}_i \cdot \vec{v}_{i2})^2}{\lambda_{i2}} + \frac{(\vec{r}_i \cdot \vec{v}_{i3})^2}{\lambda_{i3}} \right) + \lambda \left((x - x_{orig})^2 + (y - y_{orig})^2 + (z - z_{orig})^2 \right) \tag{5.27}$$

The regularization prevents instability in edge cases such as nearly-straight lines by penalizing fit vertex locations $\{x, y, z\}$ that stray from the original candidate location $\{x_{orig}, y_{orig}, z_{orig}\}$, controlled by a strength λ. By weighting the test statistic by the inverse variance along the corresponding direction, the fit preferences minimizing dimensions with a more precise measurement and therefore less variance. Furthermore, by weighting the contribution of each track segment by the variance along the trajectory, proximity to longer and more linear tracks is emphasized in the fit.

5.5.3 Track-Shower Separation

Showers are identified using three techniques, spanning a range of energies and topologies. At low energies, Multiple Coulomb Scattering (MCS) [17, 126] can be used identify electrons from muons, and nearby isolated segments of charge can indicate Bremsstrahlung photons from an electron. At higher energies, electrons produce complex showers that expand outward in a cone, and so the cone width can be used to identify them.

MCS is a process driven by numerous Coulomb scattering interactions between an ionizing particle and the surrounding medium. Since electrons are far lighter than muons, they will deflect more strongly. The numerous interactions with atomic electrons cause ionizing

electrons to travel in wiggled tracks, as opposed to the mostly straight tracks of muons. These cases are compared algorithmically by taking a track and constructing 10 cm segments from it. Then, the path length (~ 10 cm) is compared to the direct end-to-end length (≤ 10 cm) for each segment. Muon tracks measure direct lengths close to 10 cm, while electron tracks will be significantly shorter in direct length. For low energy electrons, individual Bremsstrahlung photons can produce isolated segments of charge near the electron track, providing a second method of identification. To confidently associate these segments with the candidate electron track, the isolated segment must be aligned with the beginning of the track, requiring an accurate vertex location.

High energy electrons produce showers involving an exponentially increasing number of photons and secondary electrons and positrons resulting from pair production. This causes the shower to grow wider perpendicular to the direction of its momentum, forming the shape of a cone. By comparison, muon tracks do not grow in width at all, creating a clear difference in signature. By measuring the shape and direction of the electromagnetic shower cone, the vertex of an electron can also be identified if not already known, and by measuring the perpendicular width of the cone high energy electromagnetic showers are separated from muon tracks.

5.5.4 Traditional Neutrino Vertex Identification

The traditional approach to neutrino vertex identification is developed by considering a number of observations about neutrino vertices. This leads to an overall score being computed for each vertex, and the highest score being selected. This occurs for each TPC charge cluster, and in the case that multiple clusters are matched to a single flash, one of these candidate neutrino vertices is selected, following a similar set of considerations to the individual score computations.

The most important consideration in determining the neutrino vertex is noticing that

a neutrino vertex is uniquely the originating vertex of the interaction. That means that no particles lead into it; all connected particles travel outward. By comparison, in the case of particle tracks there is always one entering particle, with all remaining tracks traveling outward. Furthermore, neutrino interactions statistically have more outgoing particles, owing to the complex nuclear interactions described in Sec. 2. Because of its high initial momentum along the beam direction, neutrino vertices are typically upstream of all other activity in the interaction.

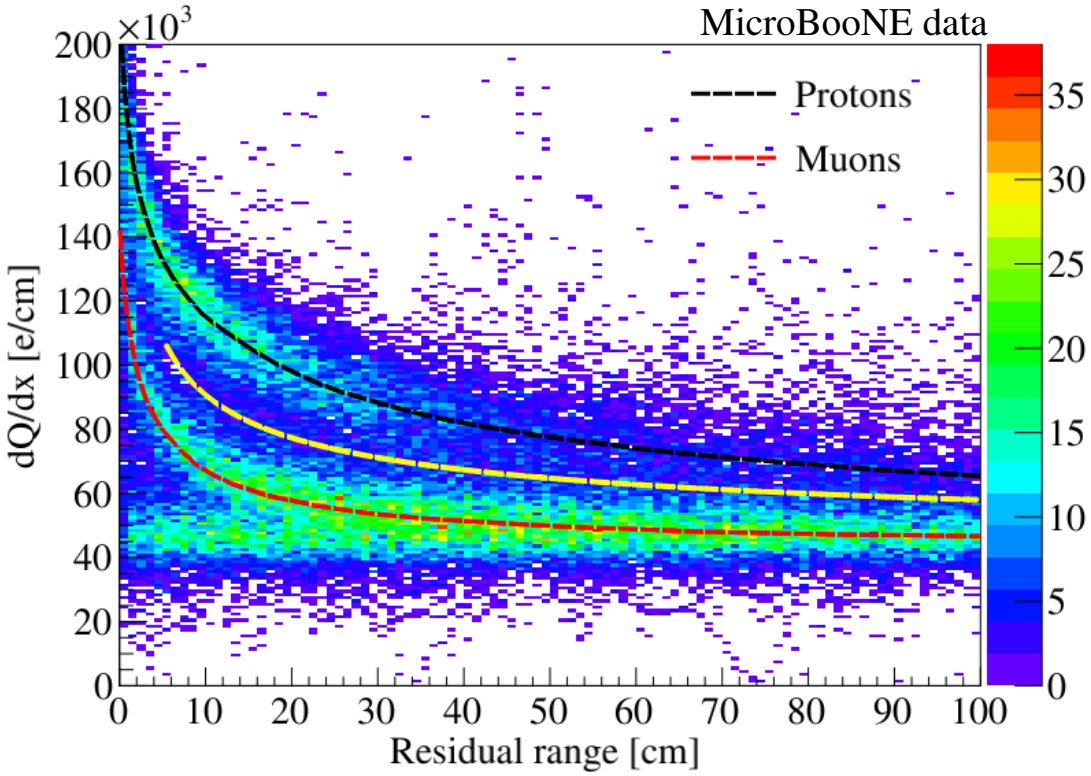

Figure 27: Separation of muon (red band) and proton (black band) tracks through the measurement of dQ/dx as a function of residual range. The yellow band indicates the selection cut between the particles. Source: [125]

The determination of particle directionality is clearly a key component of neutrino vertex identification. Shower directionality is computed following the method described previously in Sec. 5.5.3, and track directionality is determined by identifying the Bragg peak in the dQ/dx profile, as described in Sec. 5.4 and shown in Fig. 27. Limited particle identification

can also be determined using the dQ/dx profile, such as separating protons from muons and pions. Certain topologies can also be excluded such as a vertex with an entering shower and exiting track, or a vertex along a muon track resulting from a δ ray. Overall, this traditional approach is 70% accurate at determining the vertex location within 1 cm in the case of muon neutrino interactions, and this result is improved upon through the inclusion of machine-learning-based vertex finding.

5.5.5 Deep Neutral Network Vertexing

To improve the vertex identification performance, particularly on ν_eCC interactions, a Deep Neural Network (DNN) is constructed. The goal of the network falls under the category of region proposal [127], with the specific output chosen to be a prediction at each voxel of the distance to the neutrino vertex, which greatly improves training compared to a simple yes/no vertex location prediction. Following the work in other MicroBooNE analyses [128], the SparseConvNet package [129, 130] is used because it is designed for sparse datasets like a LArTPC readout, and in particular Sparse U-Net [131] is used to extract a feature vector at each voxel. The reconstructed 3D points are voxelized using 0.5 cm cubes. An illustrative diagram of the Sparse U-Net structure is shown in Fig. 28, which shows a 2-level implementation. Levels are added by iteratively inserting more "Concat-Join" blocks, and the Wire-Cell implementation uses 5 levels.

The network takes in position and charge information tensors for the voxels and produces a confidence value for each to form a Confidence Map [132], describing the distance from the neutrino vertex. The truth labels C_{true} are computed before voxelization, using the neutrino position $\vec{r}_{true}$ and a regularization parameter $\sigma = 1$ cm:

$$C_{true}(\vec{r}) = \exp\left(-\frac{||\vec{r} - \vec{r_{true}}||^2}{2\sigma^2}\right) \tag{5.28}$$

Fig. 29 shows charge information and C_{true} labels on a simulated neutrino interaction before

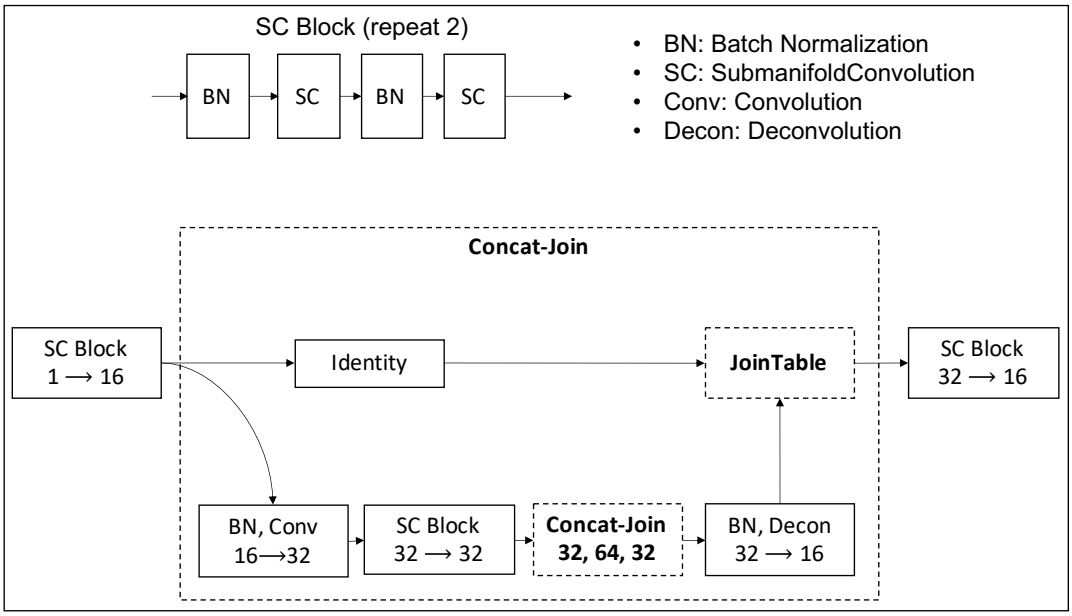

Figure 28: Illustration of a 2-level implementation the structure of Sparse U-Net. Source: [125]

voxelization.

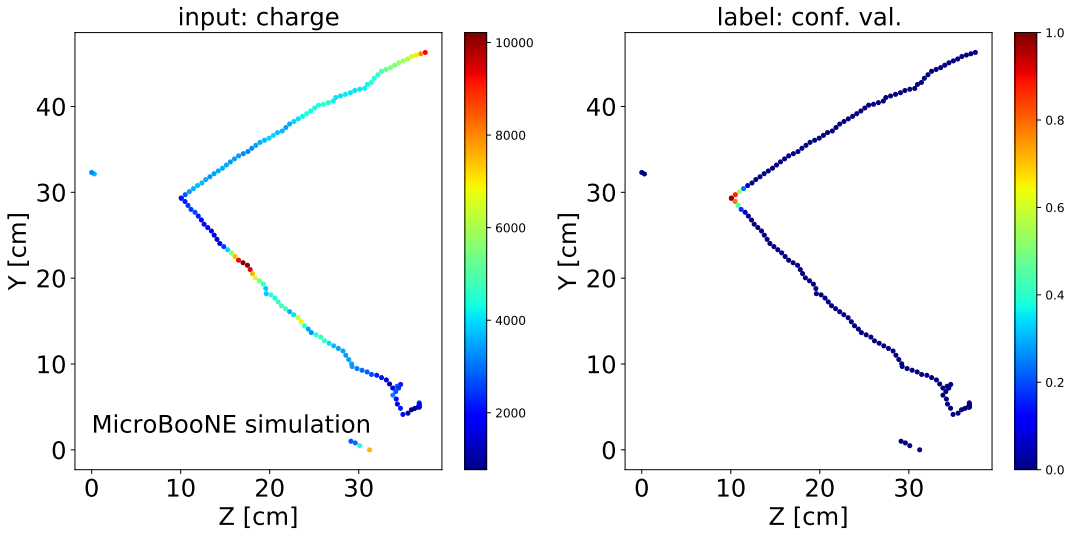

Figure 29: Example of a simulated neutrino interaction with charge info (left) and truth confidence value (right) shown. Source: [125]

The network was trained on 48k simulated ν_eCC events using the Adam optimizer [133]

with the loss function $\mathfrak{L}$:

$$\mathfrak{L} = \frac{1}{N_{voxel}} \sum^{N_{voxel}} ||C_{pred} - C_{true}||^2 \tag{5.29}$$

and an exponentially decaying learning rate to allow for fast early learning followed by precise refinement. The network was validated using 4k simulated ν_eCC events and tested on 4k simulated ν_eCC events and 4k simulated ν_μCC events. Although training focused on the ν_eCC interaction as vertexing is more difficult on it, improvement was found for both ν_eCC and ν_μCC vertexing.

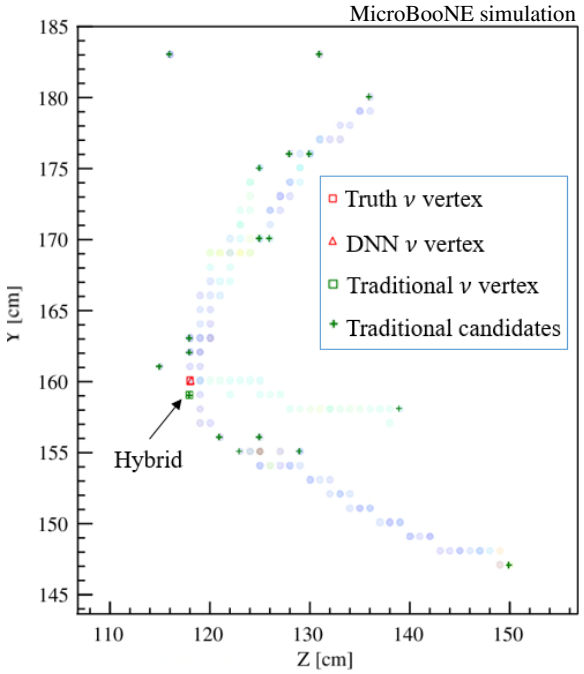

Figure 30: A simulated neutrino interaction demonstrating the performance of traditional, DNN, and hybrid vertex selection. Each filled circle represents a voxelized 3D charge projected on the Y-Z plane. Source: [125]

Existing reconstruction algorithms rely on the neutrino vertex being one of the candidate vertices found in Sec. 5.5.2. Therefore, a hybrid neutrino vertex selection algorithm is used that incorporates the DNN vertex finding information into its decision. From the best DNN neutrino vertex voxel, the 3D position is computed and the nearest vertex candidate is chosen

as the DNN vertex. If it is within 2 cm of the traditional vertex, the DNN vertex is used, otherwise the traditional vertex is kept. A visual comparison of the traditional and DNN vertex selections is shown in Fig. 30. The hybrid vertex selection leads to a 30% improvement in vertex identification within 1 cm for ν_eCC interactions and a 10% improvement for ν_μCC interactions, shown in Fig. 31.

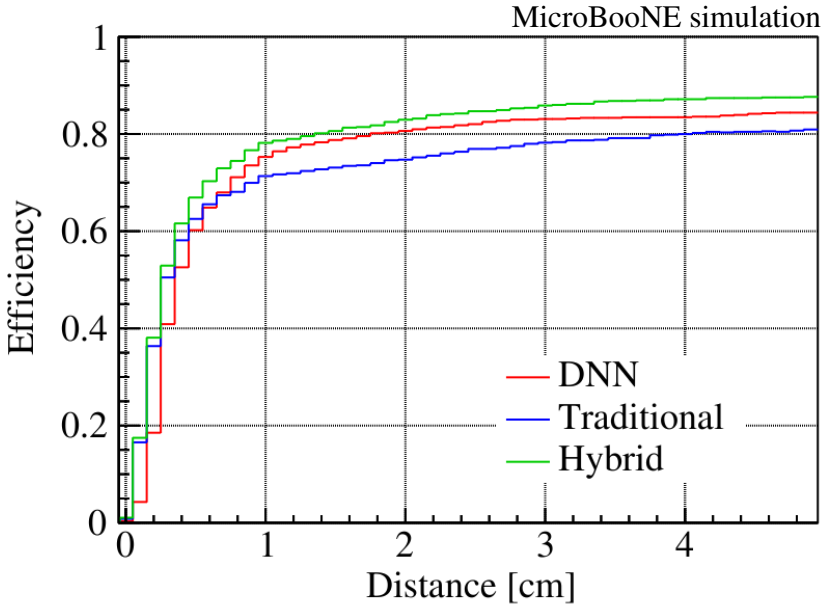

Figure 31: ν_μCC vertex identification efficiency as function of the maximum distance between the reconstructed and truth vertex position, using traditional, DNN, and hybrid vertex selection. Source: [125]

5.5.6 Particle Flow Tree

The particle flow tree is a complete, hierarchical description of the visible activity in a neutrino interaction. It identifies each particle, their dependencies (parent, children), and their energy. The steps of segmenting, trajectory fitting, vertexing, and neutrino vertex identification have already enabled a significant fraction of a particle flow tree to be constructed. What remains is the complete the categorization and organization of disconnected electromagnetic showers.

The previous shower identification algorithm struggles to fully capture the numerous disconnected segments in an electromagnetic shower. Since photons are invisible to the

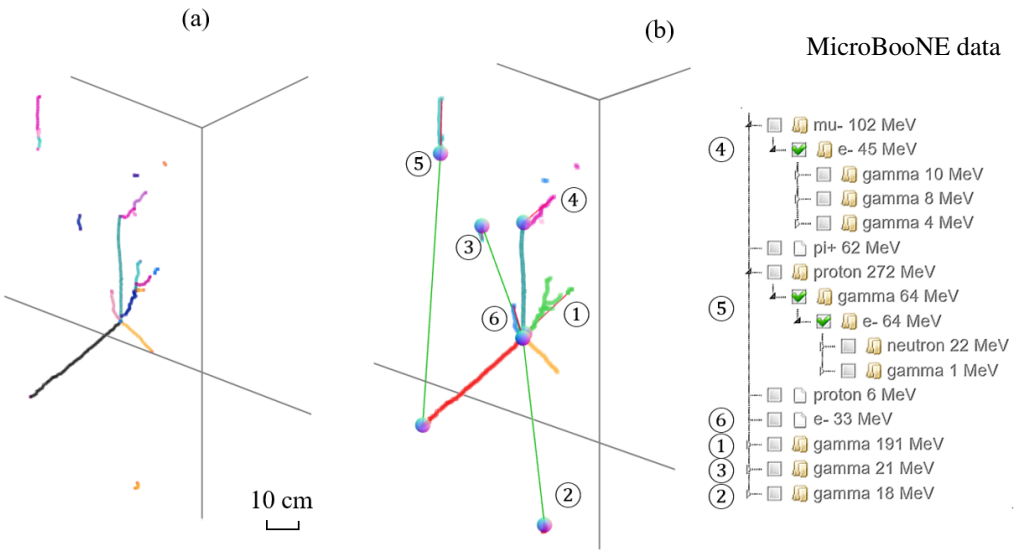

Figure 32: Illustration of electromagnetic shower grouping and the overall particle flow tree. (a) shows each track segment in separate colors, and (b) shows the same after grouping the electromagnetic shower and introducing the particle flow tree. In the particle flow, green lines show connections between spatially separated particles, and red lines connect the start and endpoints of showers. Source: [125]

LArTPC, showers appear as a large number of disconnected electron segments. To group all these segments, a 15° wide cone is extended 80 cm along the shower direction, with any isolated segments in the cone being assigned to the shower. Then, remaining isolated segments are examined individual, and potentially assigned to nearby showers based on proximity. An example of grouping the segments of an electromagnetic shower is shown in Fig. 32

With electromagnetic showers fully identified, it is possible to reconstruct the decay of invisible π^0 particles, which 99% of the time follows $\pi^0 \rightarrow \gamma\gamma$. Reconstruction of π^0 is useful for validating the electromagnetic shower energy scale, identifying the interaction type for NC and exclusive channel analyses, and to accurately reconstruct the neutrino energy. In the case of NC interactions, the neutrino vertex may have no visible directly connected particles, causing it to be misidentified, such as at the location of an electromagnetic shower. If there are at least two reconstructed electromagnetic showers, PCA [114] is used to determine the

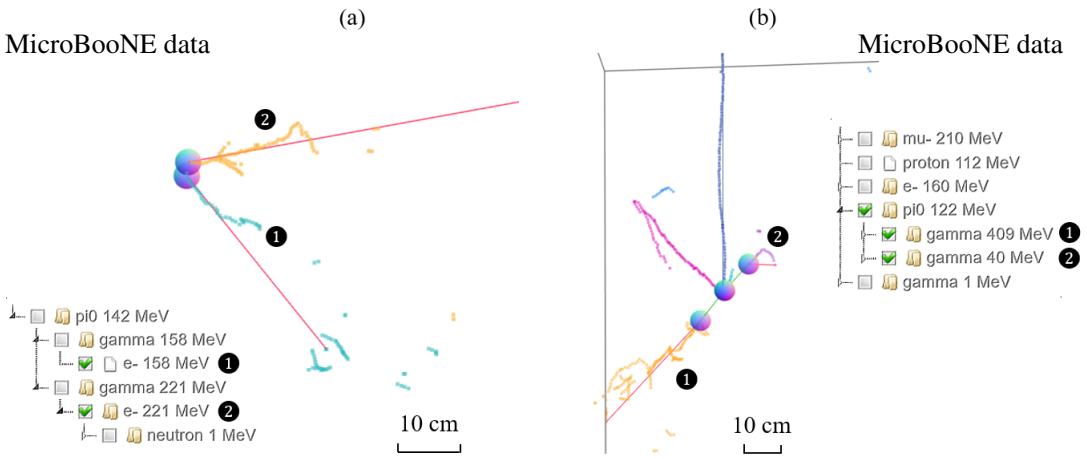

Figure 33: Examples of reconstructed π^0. In (a) a 142 MeV/c^2 π^0 is reconstructed without any TPC activity at the reconstructed vertex, and in (b) a 122 MeV/c^2 π^0 is reconstructed from the primary neutrino vertex. Photons are labeled in each event, and distinct segments are colored separately for identification. Source: [125]

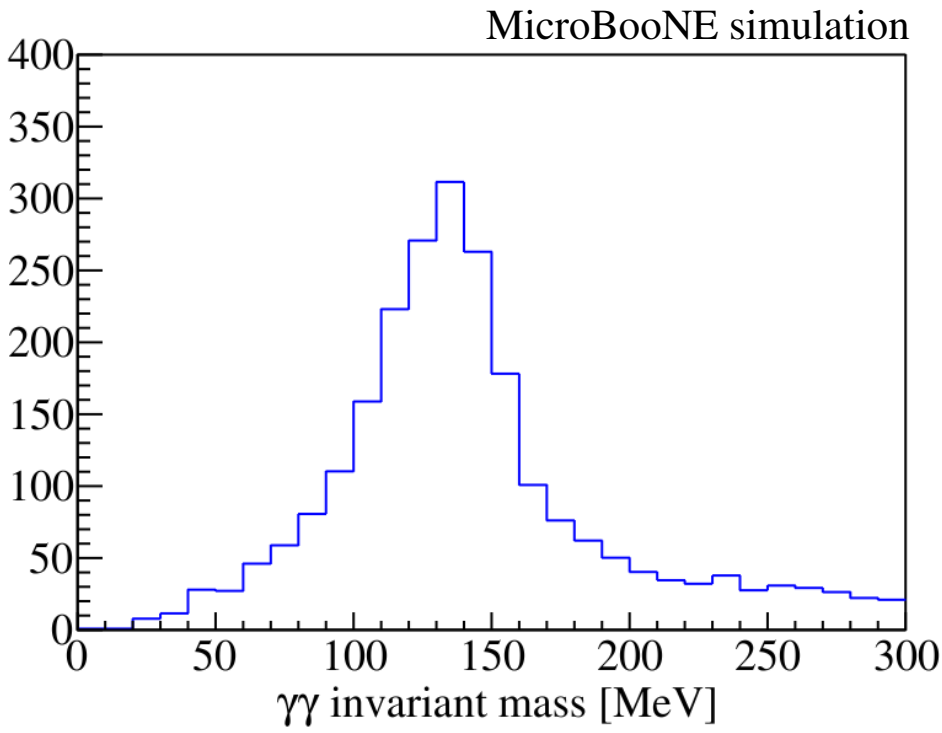

Figure 34: Reconstructed $\gamma\gamma$ invariant mass on simulated ν_μCC interactions. The tail at the high reconstructed mass is the result of incorrect association of γ candidates. Source: [125]

primary axis of each. The closest approach of each pair of primary axes is located, consisting of the closest point on each axis. If the distance between them is small enough, a candidate π^0 is considered. The π^0 is reconstructed from the candidate with the largest summed energy of the photons, its vertex is set as the midpoint between the closest points on each axis, and the direction of the two showers are adjusted to pass through the π^0 vertex. The particle flow is also updated, so that in appropriate cases the neutrino vertex is accurately labeled. Fig. 33 shows the reconstruction of two separate π^0s, and Fig. 34 shows the distribution of reconstructed π^0 masses in a simulated ν_μCCπ^0 selection, demonstrating a 15% improved resolution over previous MicroBooNE measurements [134].

5.5.7 Energy Reconstruction

Particle flow reconstructs the neutrino energy by summing the energies of all particles attached to the neutrino vertex. Therefore, accurate neutrino energy reconstruction requires accurate energy reconstruction of each particle type. In general, once a particle has been identified using the methods in previous sections, its energy is reconstructed in one of three ways: two for different cases of tracks and one for showers. In addition to the kinetic energy calculated from the visible charge in the detector, particle masses are added to the energy calculation in the case of muons and pions, and a nucleon binding energy of 8.6 MeV is added for each proton that was knocked out of the argon nucleus.

The majority of tracks are reconstructed from their measured travel range, using the NIST PSTAR database [135] to convert to total energy for a given travel range and particle mass. At very short travel distances, below 4 cm, the uncertainty on track length makes this an uncertain measurement, and so a second energy reconstruction method is used. This second approach sums the dE/dx measurements along a track to estimate the total energy, converting from dQ/dx to dE/dx using the modified box recombination model discussed in Sec. 4.3.2. However, this method is still not perfect, and underestimates the track energy

by ~ 10%. This second approach is also used in the case of very long tracks where δ rays are visible and bias the range-based estimate.

Electromagnetic showers have too complex a topology for either the range-based or dE/dx-based approaches to work. Instead, they have their energy reconstructed by summing all the measured charge in the shower and scaling it by a factor of 2.5 to account for the recombination effect (~ 50% loss) as well as bias in the charge reconstruction. Then, the scaled charge is converted to an energy measurement by multiplying by 23.6 eV per ionization electron. The performance of these energy reconstruction approaches will be studied in more detail in Sec. 7.

6 Wire-Cell Event Selection

As a surface detector, MicroBooNE is bombarded by cosmic ray muons that present a significant source of background. There are roughly $30-40$ cosmic rays per beam window, and only 1 neutrino per 600 beam spills, creating a default neutrino : background rate of $1:20,000$. These cosmic rays can appear similar to ν_μCC interactions, and both consist primarily of a muon track. This presents a threat to the neutrino purity of any selection, or through clumsily removing cosmic rays, the neutrino selection efficiency. A reduced efficiency and purity can harm an analysis, either through reduced statistics or through increased uncertainty on the makeup of a selection. This is essentially the difficulty MiniBooNE faced with its π^0 background that motivated the MicroBooNE experiment in the first place. Therefore, the Wire-Cell reconstruction was designed to leverage as much detector information as possible to enable a high-purity, high-efficiency selection. First a series of traditional cosmic ray tagging algorithms are employed, discussed in Sec. 6.1, already achieving over 80% purity and 80% efficiency in is referred to as the Generic Neutrino Detection (GND) [95, 136]. Then, the selection is refined by using a Boosted Decision Tree (BDT) in Sec. 6.2 that leverages the multitude of data products created in the reconstruction chain to reach 92% purity with 68% efficiency in the ν_μCC selection [137]. This analysis uses data from runs 1-3, constituting $\sim 6.4\times10^{20}$ POT of neutrino flux. The 1D and 3D distributions of selected events are studied in Sec. 6.3 and Sec. 6.4, respectively.

6.1 Generic Neutrino Detection

To achieve a purity of 50%, over 99.99% of cosmic rays need to be successfully identified and removed. The Wire-Cell selection relies on a number of algorithms developed throughout the reconstruction chain to reach this goal. First are the hardware and software triggers common to all MicroBooNE analyses that limit data taking to periods of beam spill and BNB-coincident light activity, respectively, as discussed earlier in Sec. 4.3.6. The neutrino : background rate of

1 : 20,000 is measured starting after the hardware trigger, and the software trigger reduces this by a factor of 100 to 1 : 210. Then flash-light matching removes non-BNB-coincident cosmic rays, reducing the window for cosmic coincidence from the 2.3 ms TPC drift time to the 1.6 μs beam spill duration, and reducing the cosmic ray background by a factor of 40. These selection cuts have only leveraged beam timing, PMT light information, and rough charge distribution (for the purposes of the predicted light pattern), without giving particular consideration to the differences in topology between neutrino interactions and cosmic rays, or to the detailed calorimetry information produced in LArTPCs. To further remove sources of background, these details will be used in combination with the beam timing information.

A key feature of cosmic rays that distinguishes them from neutrino interactions is their point of origin. Cosmic rays produce long tracks, typically originating outside the detector and passing through it. Those that pass all the way through, beginning and ending outside the fiducial volume, are called Through-Going Muons (TGMs), while those that begin outside the detector but come to rest inside it are called Stopped Muons (STMs). An easy way to remove these cosmic rays is to require that a candidate neutrino interaction be Fully Contained (FC), by which its entire ionization activity must occur within the detector fiducial volume. This contrasts with events that are Partially Contained (PC) and exist in part outside of the fiducial volume. While a fully contained selection cut would remove over 95% of remaining background events, it would also remove partially contained neutrino interactions. Especially at high energies, neutrino interactions can produce long muon tracks that exit the detector, and in total roughly 2/3 of all ν_μCC interactions are partially contained. This represents an unacceptable loss in efficiency, and so more refined selection criteria must be selected.

The first cosmic removal algorithm is inspired by the observation that singular tracks touching the detector boundary at multiple points represent TGMs and not neutrino interactions. It works by looking at the extreme points on a cluster in $\{x, y, z\}$, as well as along the primary axis as determined by using PCA [114]. If two or more extreme points are located

at the effective detector boundary, following the mapping described in Sec. 4.3.4, the event is labeled a possible TGM. To protect against removing neutrino interactions with multiple exiting tracks, a kink-finding algorithm is run from one endpoint to another to search for kinks that indicate a more complicated topology than a simple cosmic ray muon. If kinks are detected, the event is no longer considered a candidate TGM. Some TGMs may avoid detection under this algorithm because of non-functional wires at the detector boundary, preventing the reconstruction of charge in the area and the subsequent identification of a boundary intersection. To prevent this, when an extreme point lies near the detector boundary, test points are extended along the principle axis out to the detector boundary. If all test points lie in regions of non-functional wires, the track is considered to extend to the boundary in truth, and is labeled as a TGM. The removal of TGMs at this stage results in a factor of 6 reduction in the overall background.

To further separate PC cosmic rays from PC neutrino interactions, directionality information from the trajectory fit algorithm in Sec. 5.4 is leveraged. Neutrino events in the signal definition will always originate inside the fiducial volume and travel outward, while most cosmic rays will originate outside the detector and travel inward. Stopped muons can be removed from the candidate neutrino selection by identifying this directionality information. First, an event is checked to see whether it contains exactly one boundary intersection, following the extreme-points-based definition in the TGM tagger. The track is checked for kinks indicative of interaction or decay vertices. If there is no vertex, or one vertex located near the stopping point, the event is considered a candidate STM (as muons can decay to Michel electrons [138]); otherwise it is considered a candidate neutrino interaction. For candidate STMs, the trajectory is fit up to the Michel electron vertex if it exists, or endpoint otherwise, to allow dQ/dx to be measured. The dQ/dx along the final 35 cm of the candidate muon track is compared using a KS test [119] against two hypotheses: a STM profile ending in a Bragg peak and a MIP profile with a flat dQ/dx distribution. For both predictions, the

PSTAR [135] database is used to generate the dE/dx profile, which is converted to dQ/dx using the recombination model described in Sec. 4.3.2. Based on the difference in KS test scores, the event is either considered an STM and removed, or kept as a candidate neutrino interaction. The removal of STMs at this stage results in a further factor of 3 reduction in the overall background. Fig. 35 shows the detection of a cosmic ray using the STM tagger.

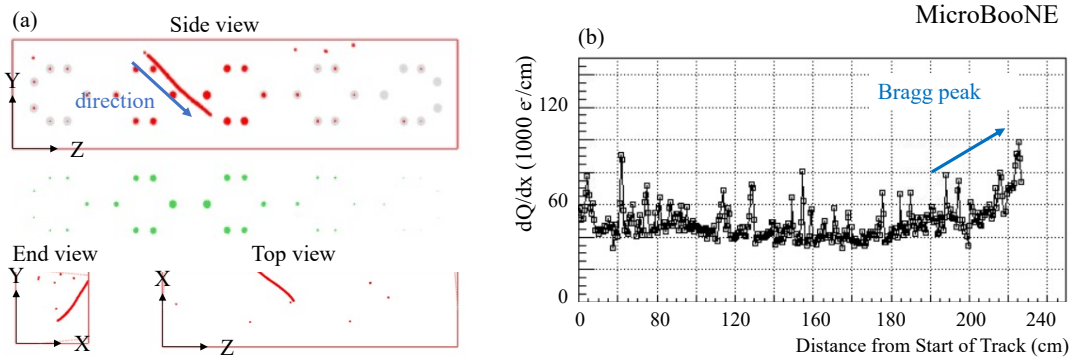

Figure 35: An example of a stopped muon from MicroBooNE data. The boundary intersection can be seen in the end and top views. Source: [95]

Remaining background events consist primarily of non-BNB-coincident cosmic rays that were incorrectly matched to the neutrino flash. Many Light Mismatch (LMM) events are the result of either small or cathode-side charge clusters that lead to a small measured flash that is difficult to correctly match, or of inaccuracies in the light predictions generated by the photon library. Although the true interaction time of the matched cosmic ray lies outside the beam spill window, this incorrect match allows them to survive the flash-light matching selection cut. Furthermore, an incorrect timing gives them a false x-axis position within a drift window, hiding their detector boundary intersections and allowing them to avoid detection through the STM and TGM taggers. These events are detected by reconsidering candidate neutrino flash-charge match quality in conjunction with a search over the other measured flashes.

First, the existing flash match KS test score is examined, and very low scores are immediately tagged as LMM events. If the KS test score is moderate, the matched cluster is

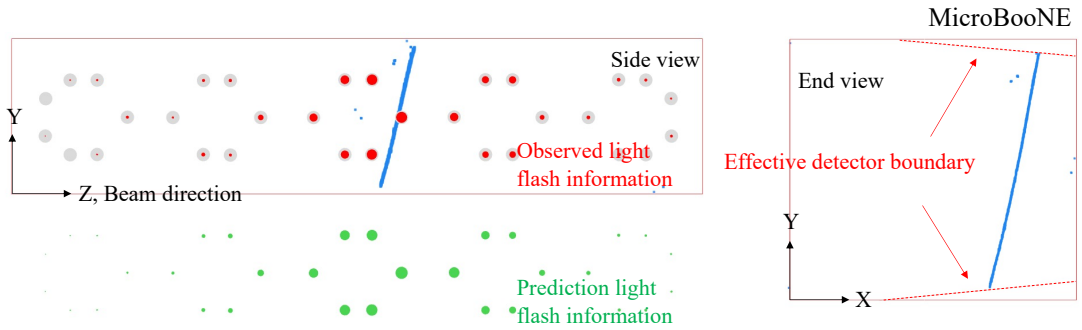

Figure 36: Example of a TGM crossing throughout he TPC at the effective detector boundary as a result of space charge distortions. The red (green) circles represent the measured (predicted) light flash PE. Source: [95]

tested against other flashes, again using a KS test between predicted and measured light. For any potential match, the new boundary intersections under the matched drift time are considered. If there are one or two boundary intersections, the cluster is evaluated with the corresponding STM or TGM tagger, respectively, and removed if found to be a cosmic ray. Additionally, even in the absence of a flash match the cluster is translated along the drift axis until it intersects the detector boundary. If multiple intersections occur simultaneously under a strict tolerance (to prevent mislabeling), the requisite flash for this drift location is assumed to be lost in the light detection or flash reconstruction steps, and the cluster is evaluated under the TGM tagger and removed if found to be a cosmic ray. Fig. 36 demonstrates the case of identifying a TGM at the effective detector boundary. Note that the effective detector boundary varies with drift distance, so that intersections with the anode and cathode are subdominant under a varied drift time.

To produce a selection of simulated events with a high degree of fidelity to the selection performance on real data, a special data product called overlay events are created. First, data is taken of background, largely cosmic rays, from a time window non-coincident with the BNB spill time. The set of these background events is referred to as EXTBNB or EXT, for external-to-the-BNB. An overlay event consists of a simulated neutrino interaction overlaid onto the readout measurements of an EXT event. This way, real background measurements

Table 2: Summary of the cumulative ν_μCC selection efficiency, cumulative (relative) background reduction, and overall background rate. Errors represent statistical uncertainties only. Source [95].

Selection Cut	ν_μCC Efficiency	Background Reduction	ν : Background
Hardware Trigger	100%	1(1)	1 : 20,000
Software Trigger	$(98.31 \pm 0.03)\%$	$(0.998 \pm 0.002) \times 10^{-2}(0.01)$	1 : 210
Charge-Light Matching	$(92.1 \pm 0.01)\%$	$(2.62 \pm 0.04) \times 10^{-4}(0.026)$	1 : 6.4
TGM Rejection	$(88.8 \pm 0.01)\%$	$(4.4 \pm 0.2) \times 10^{-5}(0.17)$	1.1 : 1
STM Rejection	$(82.9 \pm 0.01)\%$	$(1.4 \pm 0.1) \times 10^{-5}(0.32)$	2.8 : 1
LMM Rejection	$(80.4 \pm 0.01)\%$	$(6.9 \pm 0.6) \times 10^{-6}(0.50)$	5.2 : 1

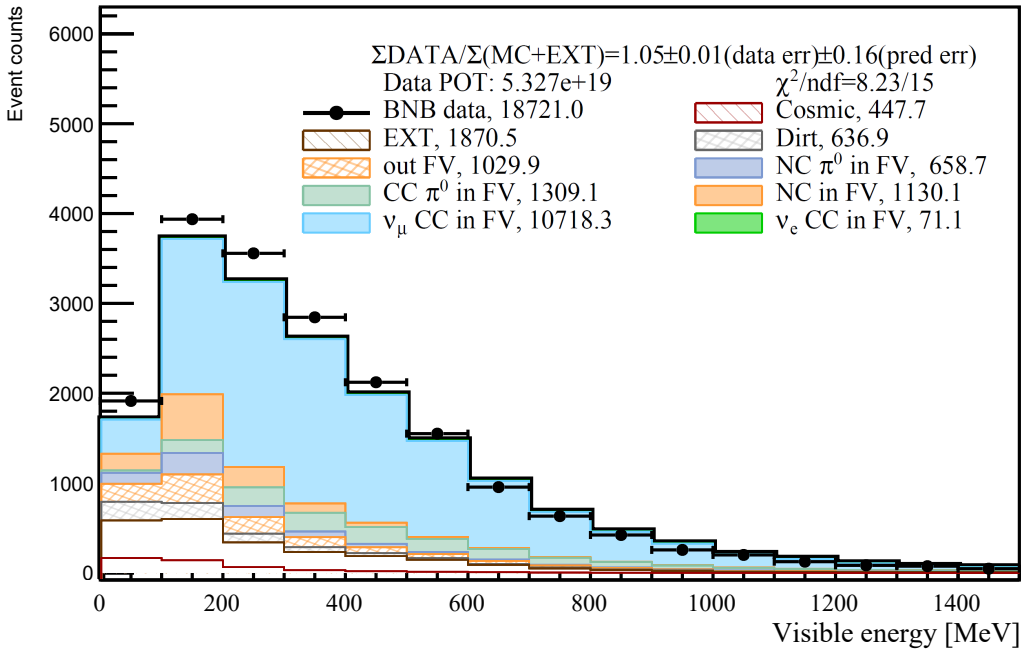

Figure 37: Distribution of selected events under generic neutrino detection over reconstructed visible energy (sum of all reconstructed charge with 2.5 scaling factor applied). The breakdown of predicted signal and backgrounds shown in each bin.

without model bias are used to generate a simulated event selection with reliable estimates of efficiency and purity. The terms “overlay" and “simulation" are largely used interchangeably within the context of estimated selection efficiencies and purities. Table 2 shows the performance of each selection cut leading up to the GND, and Fig. 37 shows the distribution of events over reconstructed visible energy.

6.2 ν_μCC Selection

A boosted decision tree allows the many correlated variables reconstructed previously to be used in combination to remove background events, even when few individually demonstrate a high selective power. A human scan of the remaining backgrounds was performed to detect their characteristics in common and salient kinematics for the purposes of background identification. These variables notably include neutrino vertex location, particle energies, and particle trajectory directionality. For example, as discussed earlier the primary muons produced in neutrino interactions rarely travel backwards or downwards. Typical backgrounds, however, include light mismatched TGMs and STMs, as well as “dirt" events, which refers to neutrino interactions outside the fiducial volume. The remaining dirt events usually feature a charged hadron that enters the detector from upstream before undergoing a hadronic interaction, which gets labeled as the neutrino vertex. In these cases the STM tagger is not well suited to identify the background, but the combined observations of a backward facing particle that is highly energetic can give a high likelihood of background identity. To remove NC events from the selection, each event is required to reconstruct a muon track at least 5 cm long attached to the vertex.

A total of over 300 such variables are used in the BDT, which is trained using the BDT package XGBoost [139]. XGBoost uses parallel tree boosting, and improves the generalizability of the trained model, allowing such a large number of variables to be used without over-fitting. The performance of the BDT is shown by comparing the ν_μCC selection efficiency and purity as a function of BDT score in Fig. 38. A selection cut is made keeping events with a score above 0.9, producing a ν_μCC event selection with 92% purity and 68% efficiency.

6.3 Event Selection Over 1D Distributions

While the ultimate goal of this analysis is to produce a triple-differential cross section measurement, single differential measurements can serve as useful stepping stones capable of

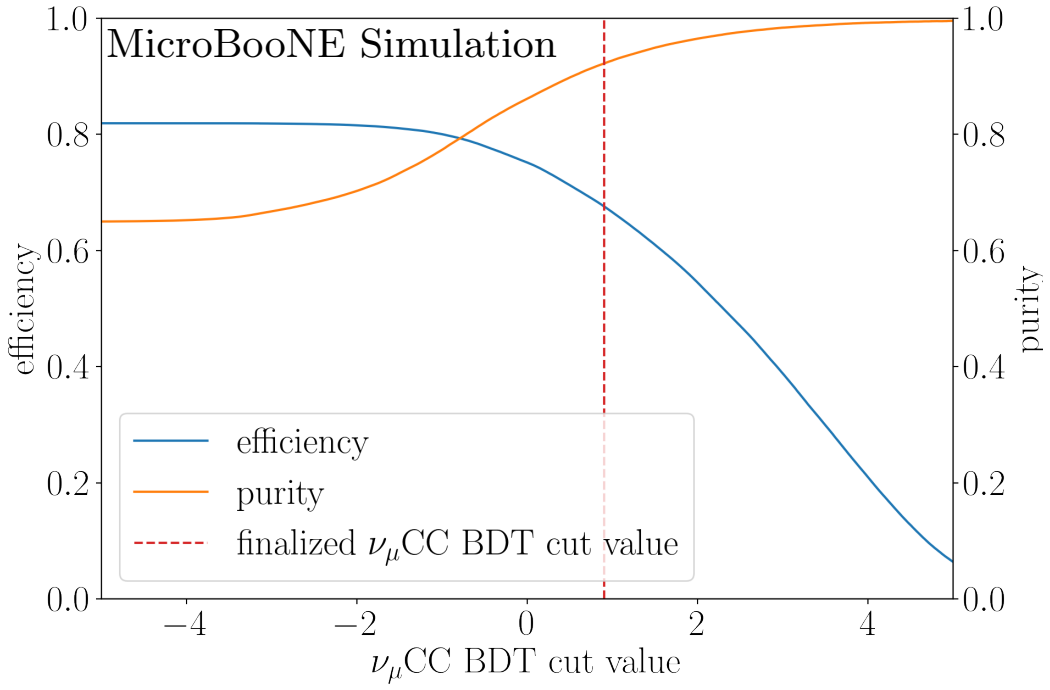

Figure 38: Performance of BDT shown by comparing ν_μCC selection efficiency and purity as a function of BDT score. The finalized event selection uses the cut at the highlighted BDT score of 0.9. Source: [137]

providing insight in their own right. Here the distribution of events is studied over muon energy E_μ, muon scattering angle $cos(\theta_\mu)$, visible hadronic energy E_{had}^{reco}, and neutrino energy E_ν. These variables are particularly interesting, as together they are capable of describing the principle neutrino-argon interaction kinematics. Since this interaction is described by thee degrees of freedom, the four reconstructed quantities are redundant, with $\{E_\mu, cos(\theta_\mu), E_{had}^{reco}\}$ sufficing. However, as was discussed in Sec. 2.1, the neutrino energy is the most physically important variable, as it drives neutrino oscillations and can be used to separate interaction channels. The neutrino energy can be expressed in terms of E_μ and the energy transferred to the argon system, denoted as ν (context is needed to tell if one means transfer energy or neutrino by ν). The transfer energy is comprised of a visible portion, measured with E_{had}^{reco}, as well as an invisible portion, $E_{had}^{missing}$, largely consisting of non-ionizing particles such as neutrons that cannot be detected in a LArTPC. Therefore, E_{had}^{reco} is the best direct measurement of

ν, and together with E_μ can be used to reconstruct E_ν. As a result, either $\{E_\mu, cos(\theta_\mu), E^{reco}_{had}\}$ or $\{E_\mu, cos(\theta_\mu), E_\nu\}$ can be seen as a suitable choice of variables to describe the kinematic phase space.

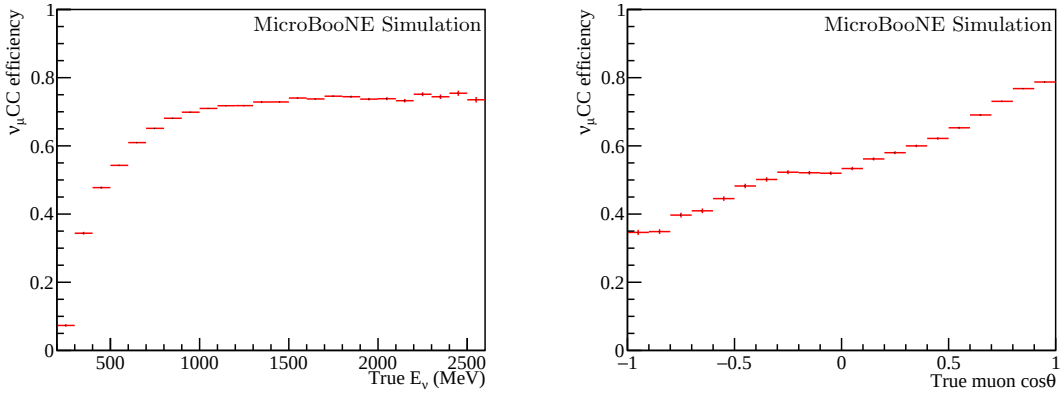

Figure 39: ν_μCC selection efficiency in simulation as a function of (left) neutrino energy and (right) muon scattering angle. Source: [137]

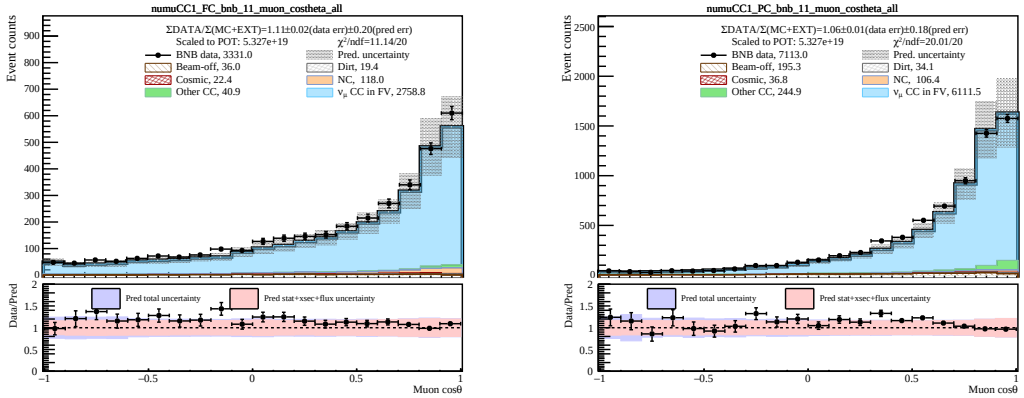

Figure 40: Distribution of selected FC (left) and PC (right) ν_μCC events over reconstructed muon angle. The predicted distribution of signal and background events are stacked in each bin. Ratio plots are shown on the bottom, with statistical uncertainties on the data measurement, and systematic uncertainties given with the blue (all uncertainties) and red (no detector uncertainties) bands.

A selection cut is applied onto the ν_μCC selection described in the previous section, requiring $E_\nu \in [0.2, 4.0]$ GeV. The reconstruction struggles to detect and reconstruct events properly below 200 MeV, and there are few events produced by the BNB above 4 GeV, making these suitable cutoffs to the selection with only 1% efficiency loss. Fig. 39 shows the ν_μCC

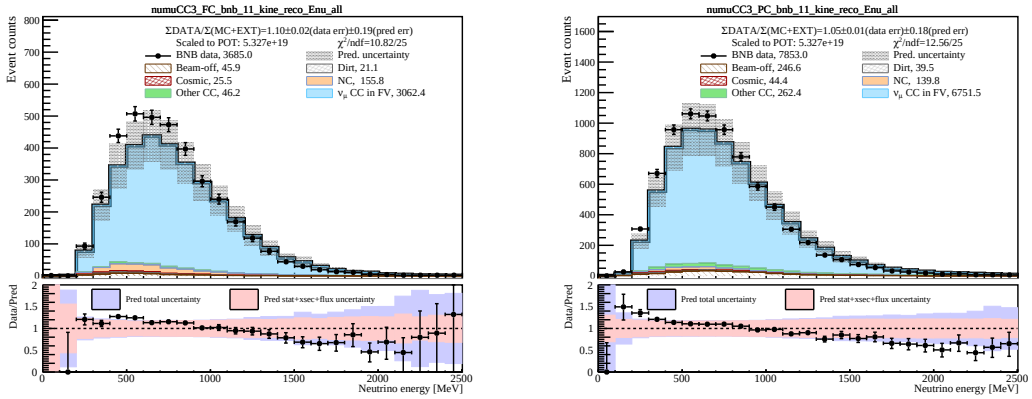

Figure 41: Distribution of selected FC (left) and PC (right) ν_μCC events over reconstructed neutrino energy. The predicted distribution of signal and background events are stacked in each bin. Ratio plots are shown on the bottom, with statistical uncertainties on the data measurement, and systematic uncertainties given with the blue (all uncertainties) and red (no detector uncertainties) bands.

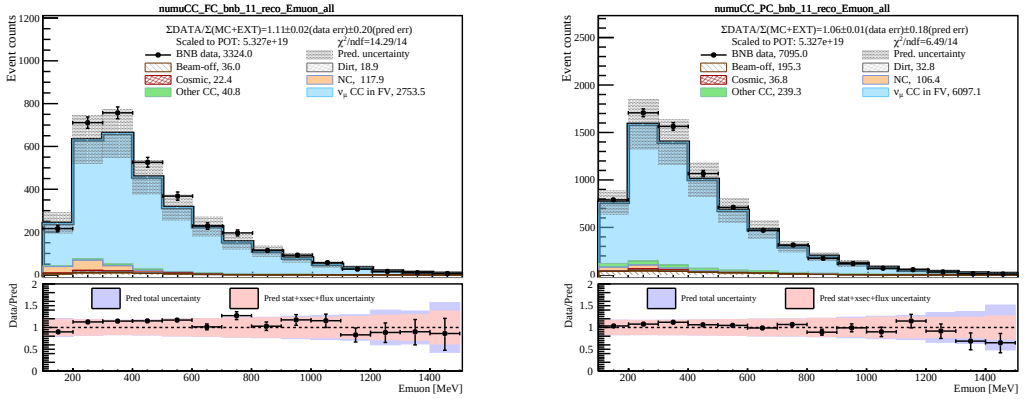

Figure 42: Distribution of selected FC (left) and PC (right) ν_μCC events over reconstructed muon energy. The predicted distribution of signal and background events are stacked in each bin. Ratio plots are shown on the bottom, with statistical uncertainties on the data measurement, and systematic uncertainties given with the blue (all uncertainties) and red (no detector uncertainties) bands.

selection efficiency in simulation over E_ν and $cos(\theta_\mu)$, showing the improved selection efficiency at high energy and forward angles. Fig. 40, Fig. 41, Fig. 42, and Fig. 43 show the signal distribution and breakdown of background sources as a function of $cos(\theta_\mu)$, E_ν, E_μ, and E_{had}, respectively, separated by FC and PC channels.

Maintaining a high selection efficiency across the signal phase space is particularly important for two reasons. First, it reduces the change that the final event distribution is biased

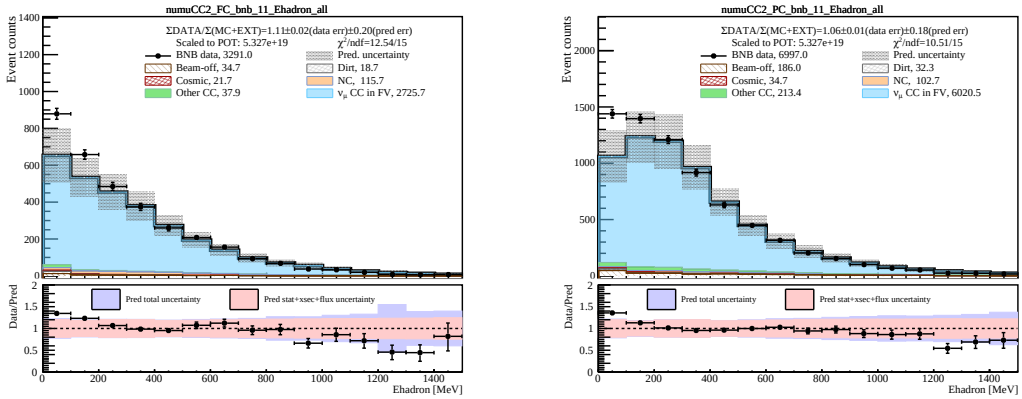

Figure 43: Distribution of selected FC (left) and PC (right) ν_μCC events over reconstructed hadronic energy. The predicted distribution of signal and background events are stacked in each bin. Ratio plots are shown on the bottom, with statistical uncertainties on the data measurement, and systematic uncertainties given with the blue (all uncertainties) and red (no detector uncertainties) bands.

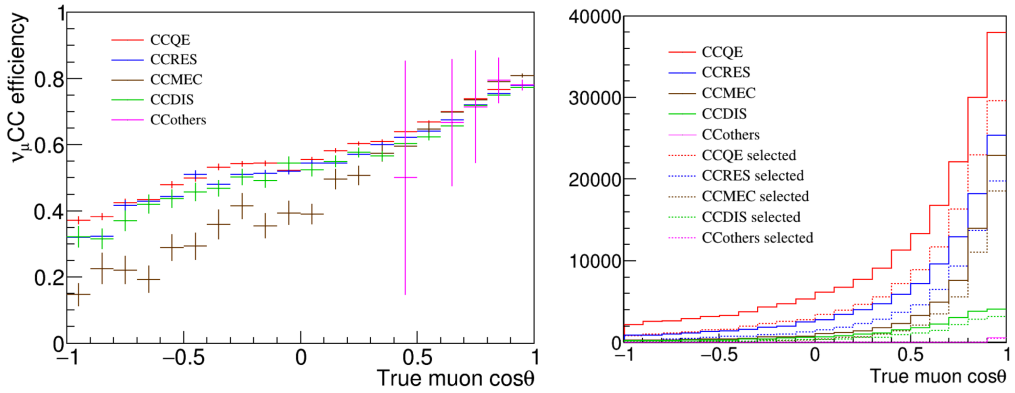

Figure 44: Simulated true muon angle (left) selection efficiency and (right) binned event distribution before/after selection.

with respect to the types of events it selects. This is checked in Fig. 44, Fig. 45, Fig. 46, and Fig. 47 in simulation over $cos(\theta_\mu)$, E_ν, E_μ, and ν, respectively. The event distributions shown on the right side of Fig. 45, Fig. 46, and Fig. 47 are grouped into the analysis binning used later to produce unfolded cross section measurements. The choice of binning is motivated by the statistics available, as well as the reconstructed variable resolutions, which will be discussed more in Sec. 7.

Each interaction channel is shown, and there are no clear holes in efficiency that would

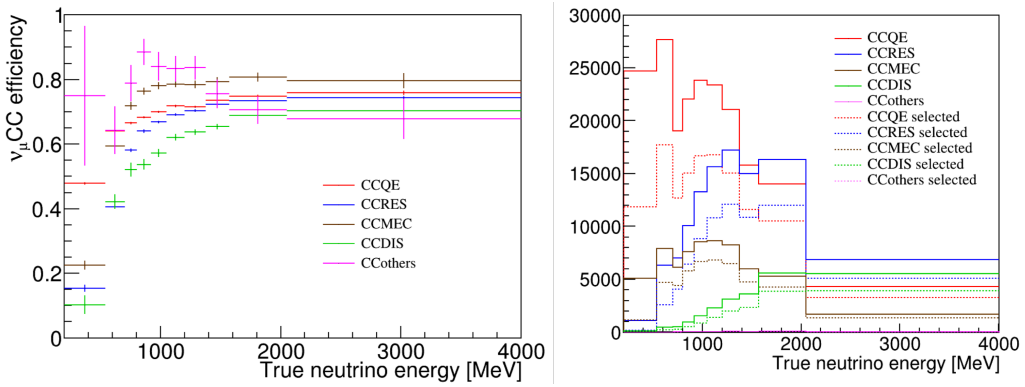

Figure 45: Simulated true neutrino energy (left) selection efficiency and (right) binned event distribution before/after selection.

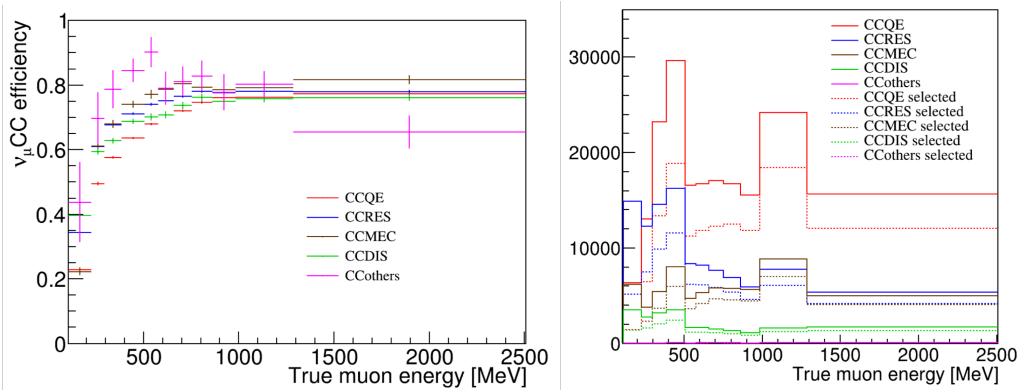

Figure 46: Simulated true muon energy (left) selection efficiency and (right) binned event distribution before/after selection.

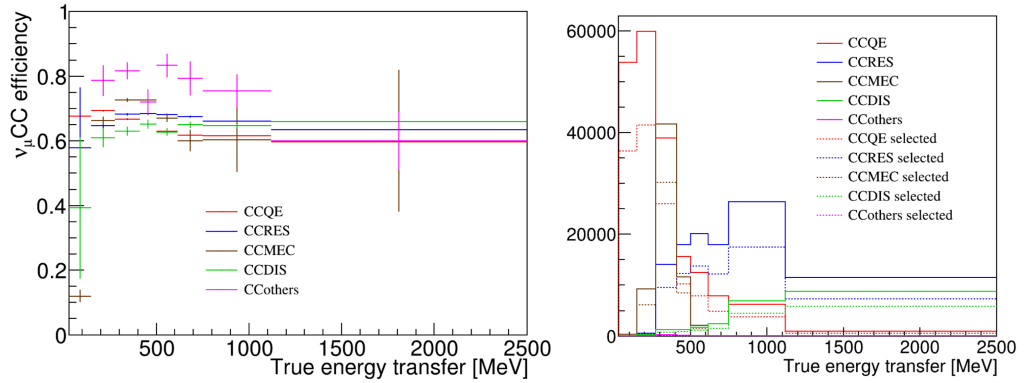

Figure 47: Simulated true transfer energy (left) selection efficiency and (right) binned event distribution before/after selection.

indicate a strong bias in the selection. In particular, the similar performance between interaction channels over muon scattering angle, which is heavily utilised in distinguishing neutrino

interactions from cosmic rays, demonstrates the inclusive nature of the selection. The other reason a selection should maintain a non-zero efficiency is so that the model prediction can be fully validated. Producing a cross section measurement requires an accurate model prediction to account for biases in the selection, such as efficiency loss. The details of model validation are discussed in detail in Sec. 9, and rely heavily on the fact that the event selection is able to test the model prediction over the analysis phase space.

6.4 Event Selection Over 3D Distribution

The triple-differential analysis uses a 3D measurement space with bins in the reconstructed variables of $\{E_\nu, P_\mu, cos(\theta_\mu)\}$, where the muon momentum P_μ is used in place of E_μ. Both variables represent the same detector information, but using P_μ allows for a more direct comparison with the previous MicroBooNE double-differential $\nu_\mu CC$ cross section measurement [140] performed over $\{P_\mu, cos(\theta_\mu)\}$. For the most accurate comparison, a final selection cut requiring the muon momentum to be less than 2.5 GeV/c is added, removing only 1.5% of events. The complete list of selection criteria is given below:

The analysis uses a reconstructed binning following a grid-based structure. Each variable is cut into a number of slices, with a 3D bin defined by its slice along each axis. The neutrino energy is reconstructed into four slices, given by the edges $E_\nu \in \{0.2, 0.57, 1.05, 1.57, 4.0\}$ GeV. The muon momentum is reconstructed in 15 slices, 0.1 GeV/c each, up to 1.5 GeV/c, plus an overflow slice up to 2.5 GeV/c. The muon scattering angle is reconstructed into nine slices, given by the edges $cos(\theta_\mu) \in \{-1, -0.5, 0, 0.27, 0, 45, 0, 62, 0.76, 0.86, 0.94, 1\}$. Fully and partially contained events are reconstructed separately, for a total of $4 \times 9 \times 16 \times 2 = 1152$ reconstructed bins.

The distribution of reconstructed events and predicted signal and backgrounds is shown in Figs. 48-55. Each figure shows data from a particular E_ν slice, with each of the nine sub-plots containing data from a particular θ_μ slice. The same distributions shown with a

breakdown of the predicted signal interaction channel are presented in Figs. 56-63. Below each distribution is a ratio plot comparing the measured values to their predicted counterpart in each bin, with statistical uncertainties on the data, and systematic uncertainties including (excluding) detector systematics shown in the purple (red) bands. Some bins have low statistics and therefore large statistical uncertainties, particularly at high neutrino energy. These bins do not negatively impact the overall unfolded measurement uncertainty, since the unfolding simultaneously maps from FC and PC measurements to fully inclusive analysis bins, as will be discussed in more detail in Sec. 10.

- Pass Wire-Cell Generic Neutrino Selection
  - Hardware Filter
  - Software Filter
  - Flash-Light Matching
  - Through-going Muon Rejection Rejection
  - Stopped Muon Rejection
  - Light Mismatch Rejection
- Pass Wire-Cell ν_μCC Selection
  - ν Vertex Inside Fiducial Volume
  - Primary μ Track $> 5\,\text{cm}$
  - ν_μ BDT Score Above 0.9
- $E_\nu \in [0.2, 4.0]\,\text{GeV}$
- $P_\mu < 2.5\,\text{GeV/c}$

“Cosmic" refers to cosmic rays coincident with the beam spill window, “EXT" refers to non-coincident cosmic rays, “dirt" refers to neutrino interactions outside the cryostat, “out

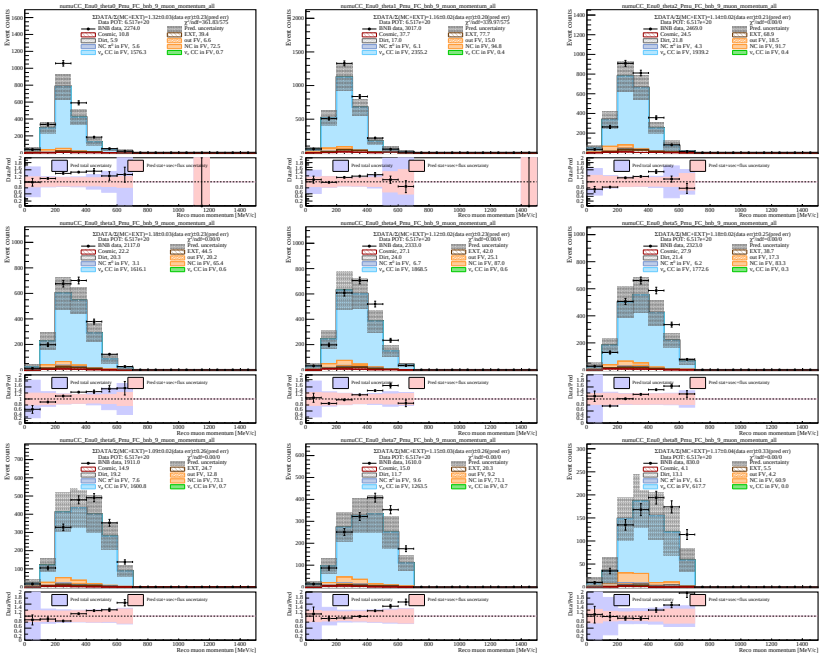

Figure 48: Distribution of selected FC events and predicted signal and backgrounds over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [0.2, 0.705]$ GeV.

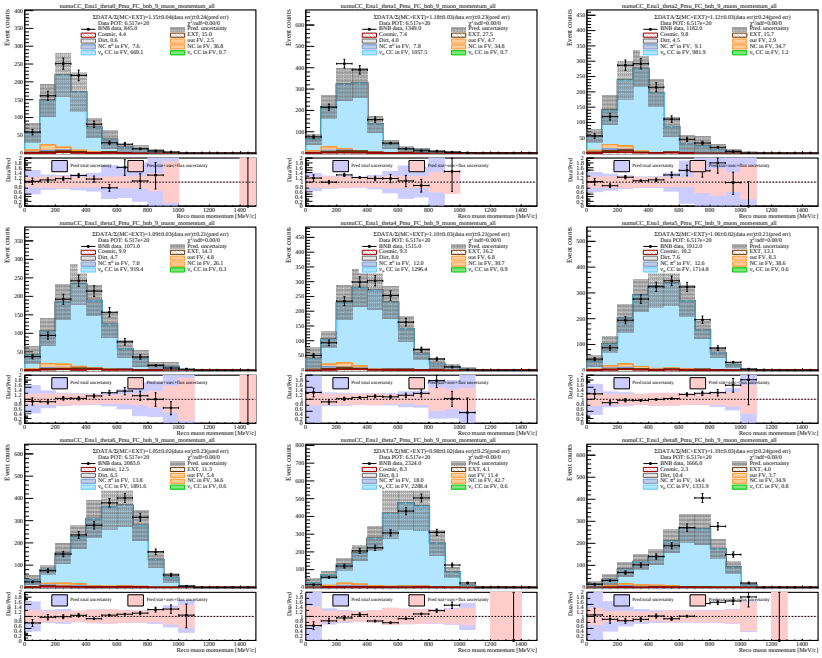

Figure 49: Distribution of selected FC events and predicted signal and backgrounds over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [0.705, 1.05]$ GeV.

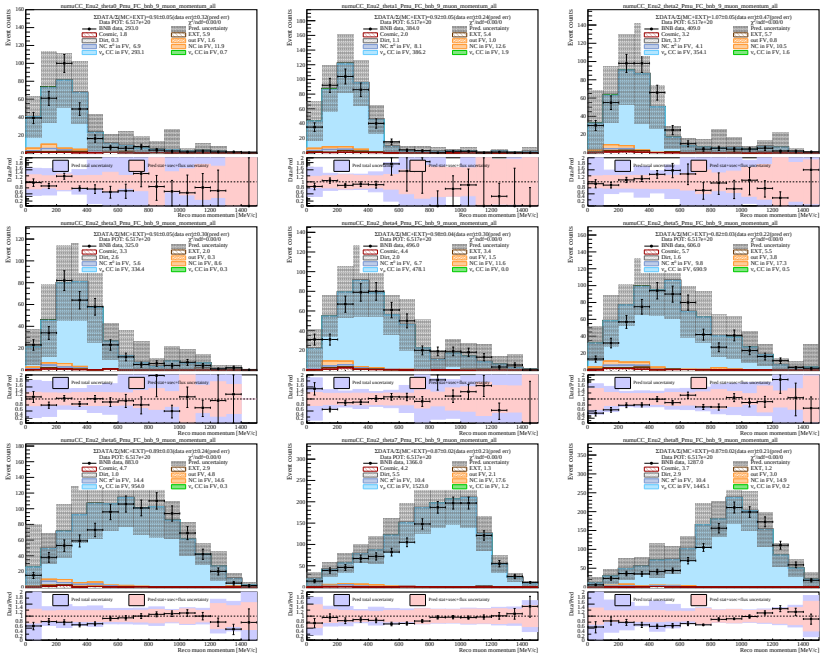

Figure 50: Distribution of selected FC events and predicted signal and backgrounds over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [1.05, 1.57]$ GeV.

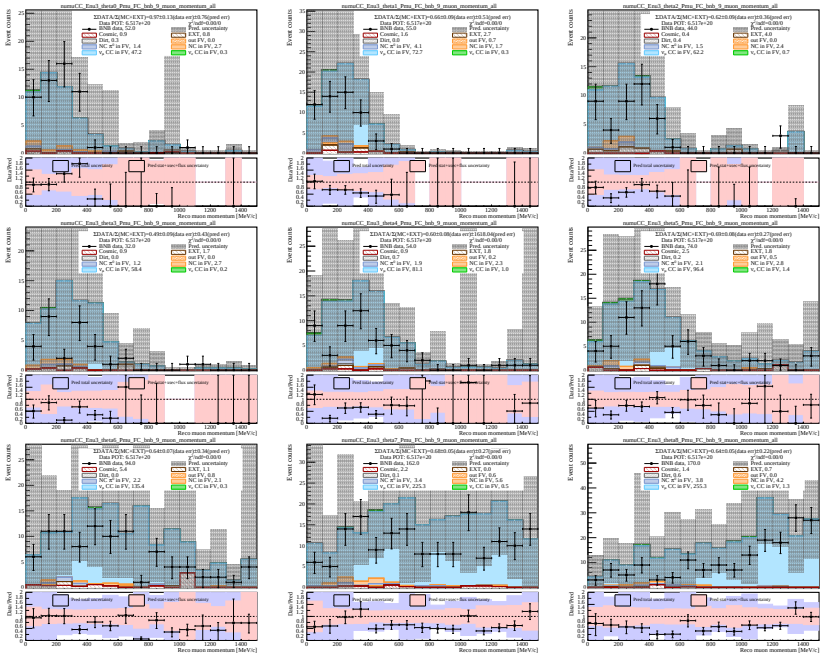

Figure 51: Distribution of selected FC events and predicted signal and backgrounds over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [1.57, 4.0]$ GeV.

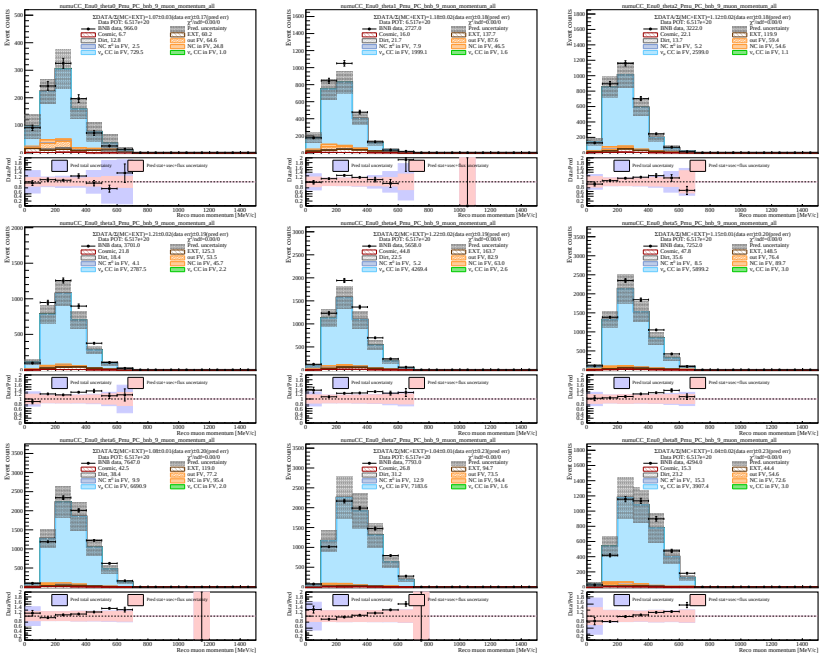

Figure 52: Distribution of selected PC events and predicted signal and backgrounds over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [0.2, 0.705]$ GeV.

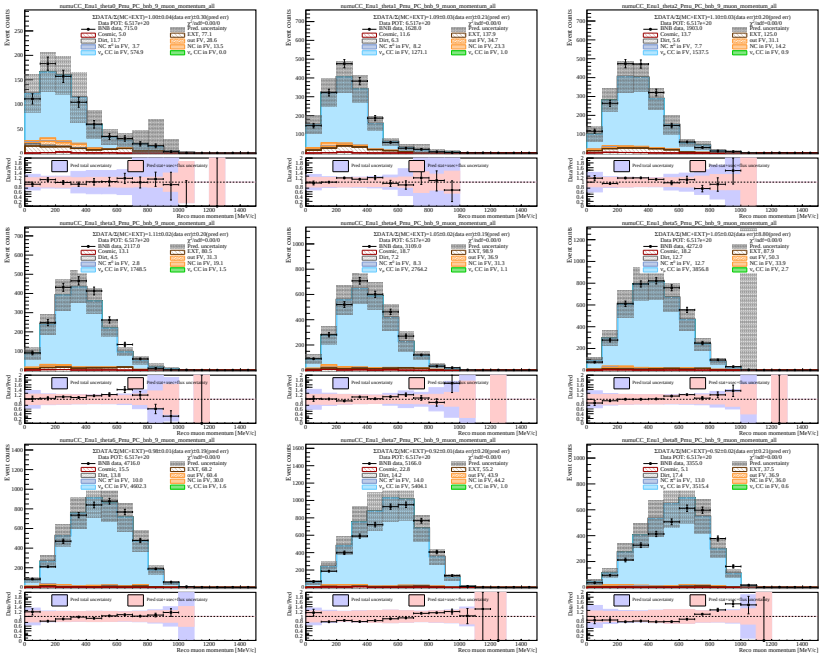

Figure 53: Distribution of selected PC events and predicted signal and backgrounds over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [0.705, 1.05]$ GeV.

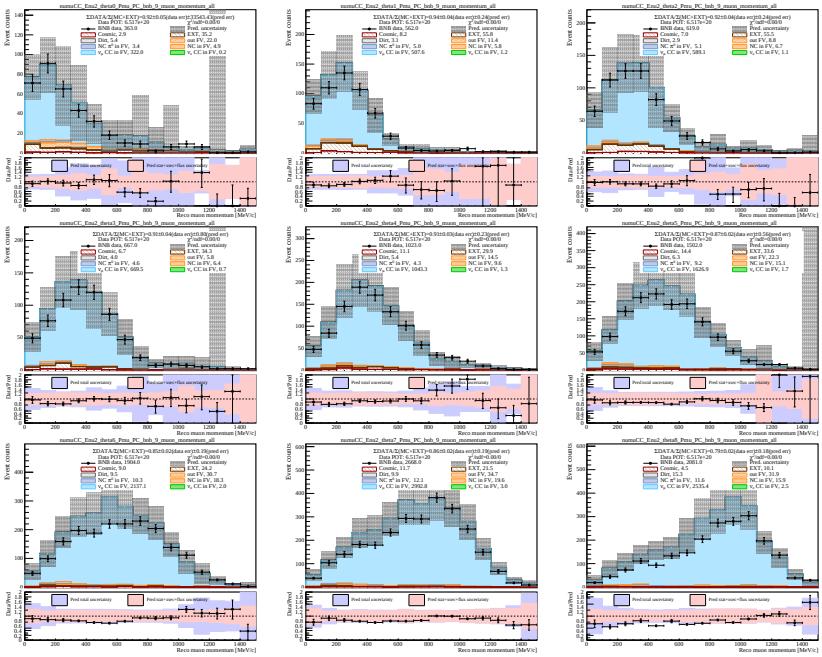

Figure 54: Distribution of selected PC events and predicted signal and backgrounds over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [1.05, 1.57]$ GeV.

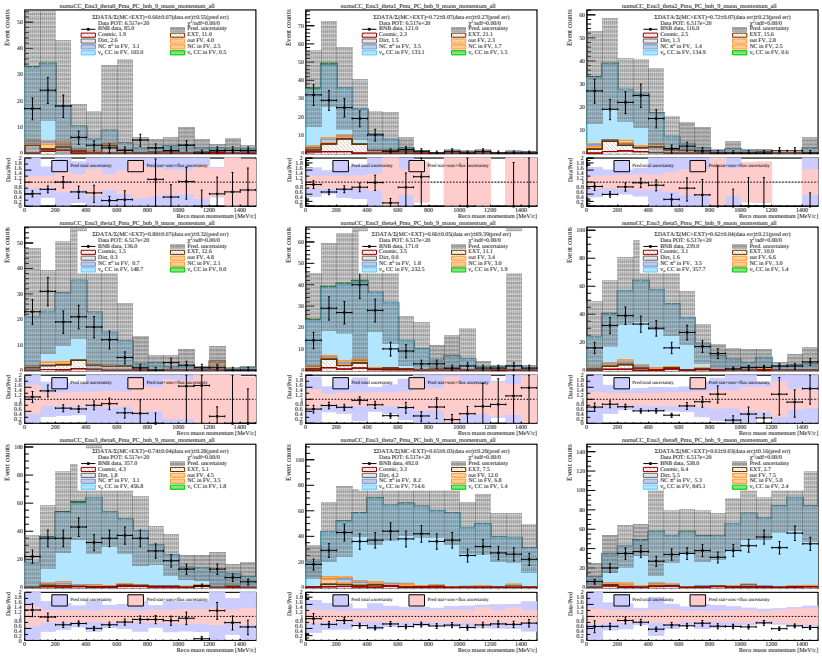

Figure 55: Distribution of selected PC events and predicted signal and backgrounds over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [1.57, 4.0]$ GeV.

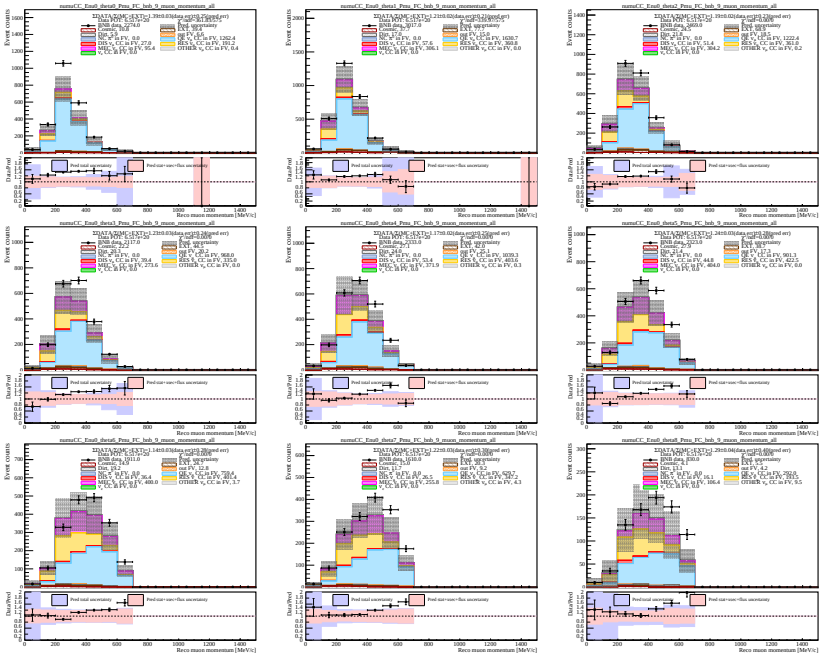

Figure 56: Distribution of selected FC events and predicted breakdown of signal interaction channel over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [0.2, 0.705]$ GeV.

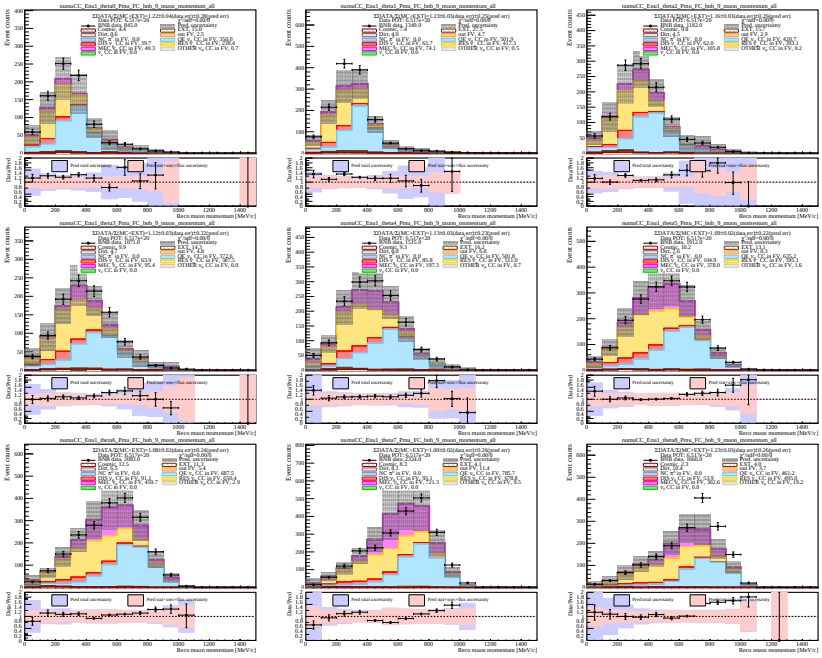

Figure 57: Distribution of selected FC events and predicted breakdown of signal interaction channel over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [0.705, 1.05]$ GeV.

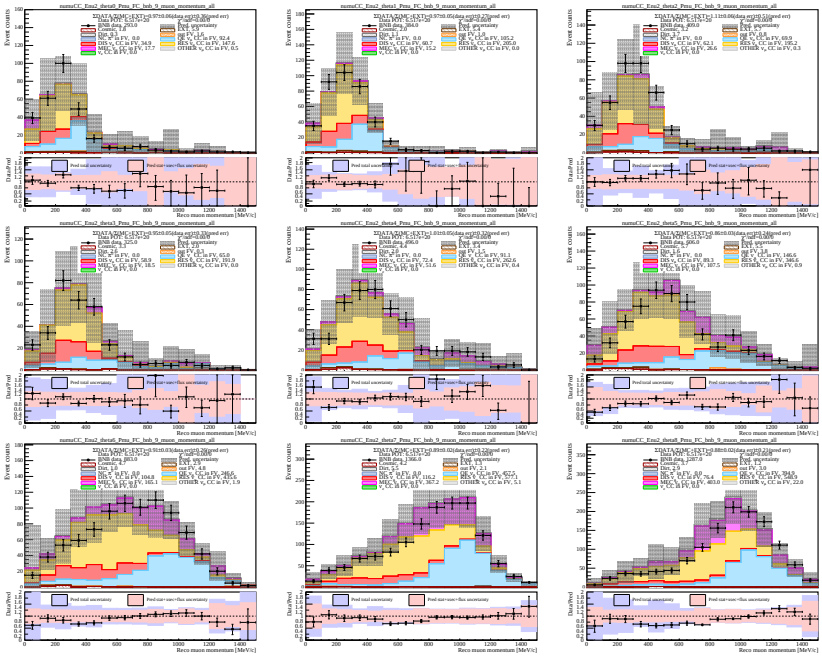

Figure 58: Distribution of selected FC events and predicted breakdown of signal interaction channel over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [1.05, 1.57]$ GeV.

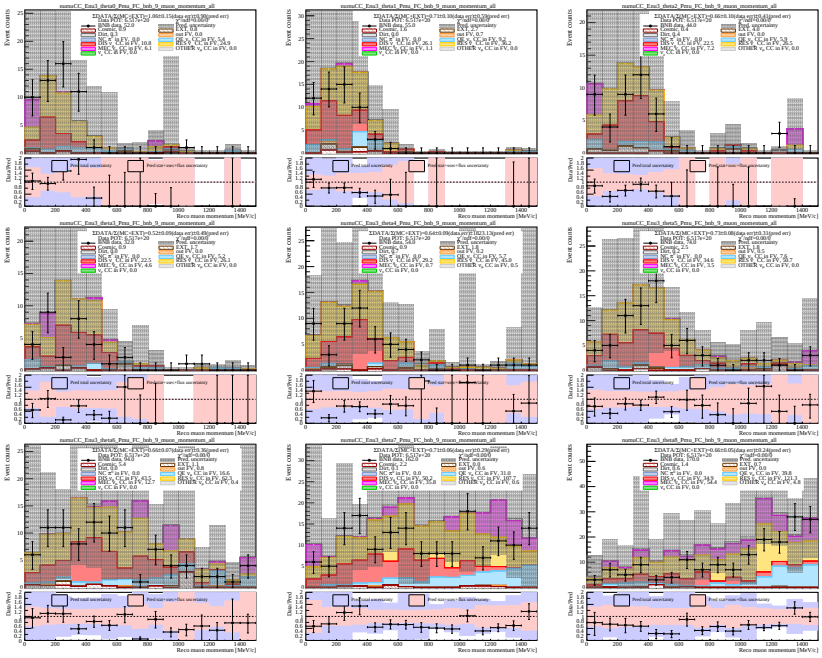

Figure 59: Distribution of selected FC events and predicted breakdown of signal interaction channel over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [1.57, 4.0]$ GeV.

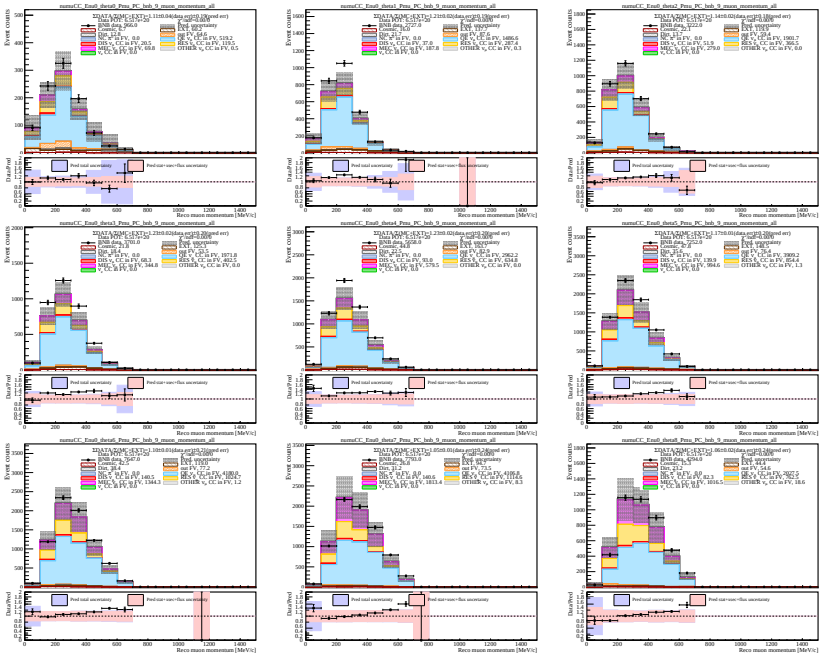

Figure 60: Distribution of selected PC events and predicted breakdown of signal interaction channel over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [0.2, 0.705]$ GeV.

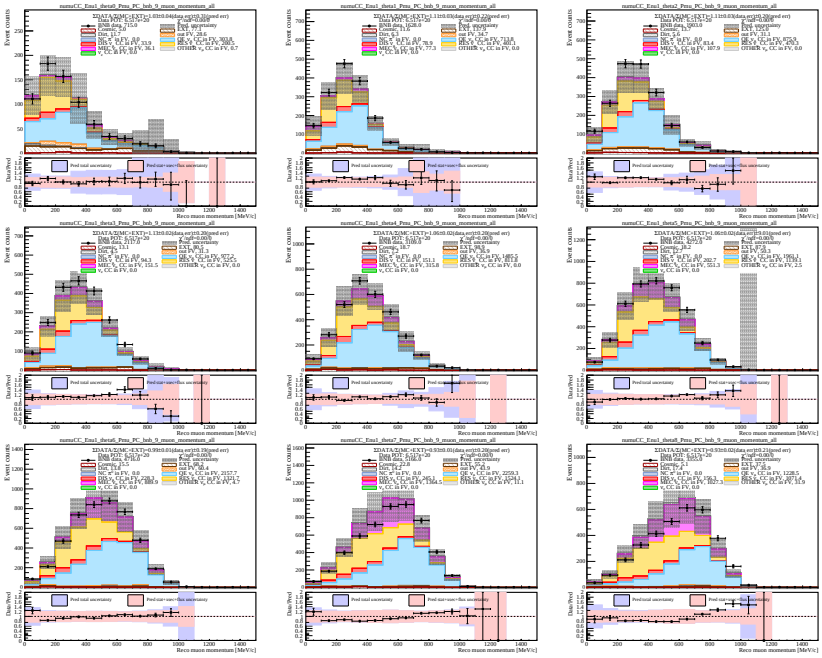

Figure 61: Distribution of selected PC events and predicted breakdown of signal interaction channel over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [0.705, 1.05]$ GeV.

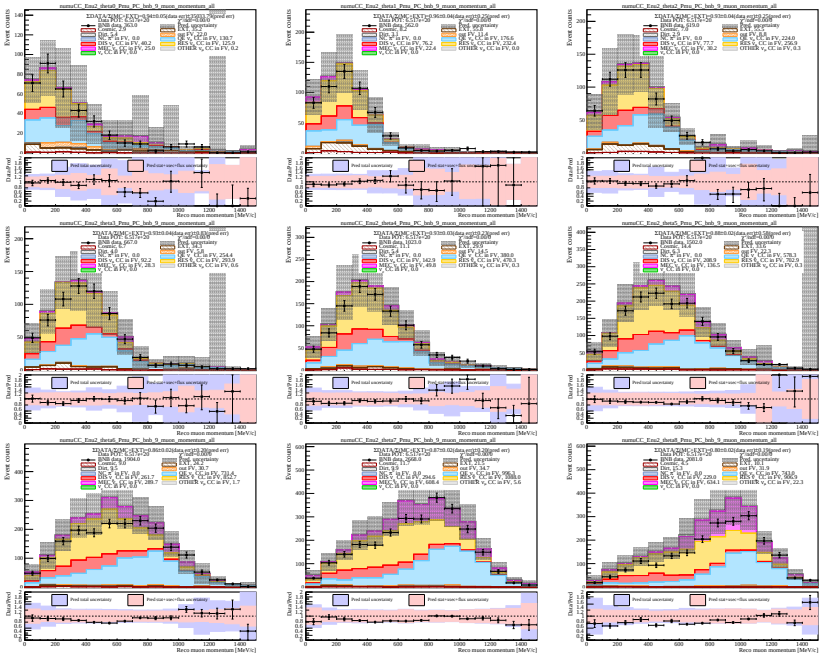

Figure 62: Distribution of selected PC events and predicted breakdown of signal interaction channel over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [1.05, 1.57]$ GeV.

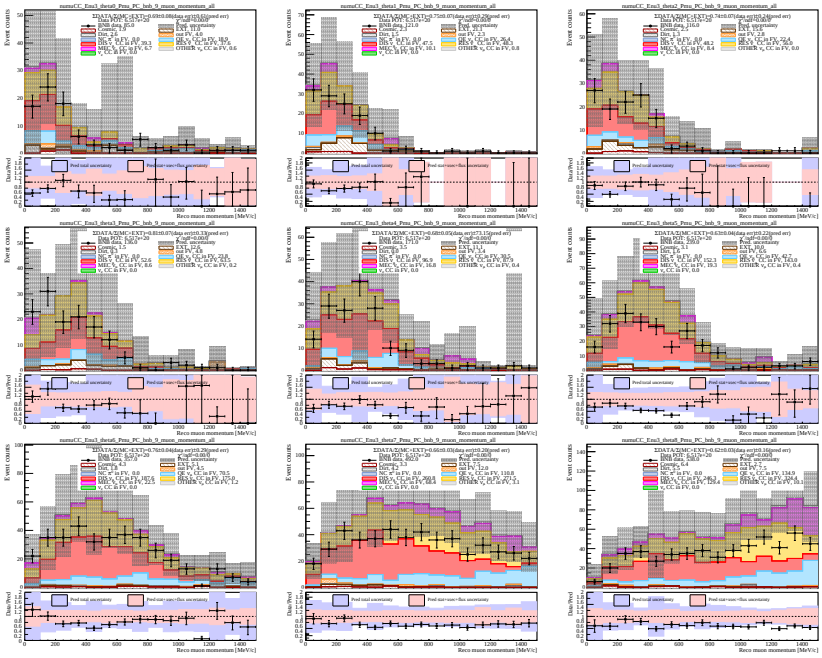

Figure 63: Distribution of selected PC events and predicted breakdown of signal interaction channel over muon momentum. Each angle slice is shown within the neutrino energy slice $E_\nu \in [1.57, 4.0]$ GeV.

FV" refers to neutrino interactions outside the fiducial volume, "NC π^0 in FV" refers to neutral current interactions with a π^0 and a neutrino vertex inside the fiducial volume, "NC in FV" refers to other neutral current interactions inside the fiducial volume, "ν_eCC in FV" refers to electron neutrino interactions inside the fiducial volume, and "ν_μCC in FV" refers to the signal channel.

Some features of the data are easily identified over the 3D distribution. There is a clear shift in the distribution of muon momentum across energy ranges and angles, with higher energies and more forward angles experiencing a higher peak momentum. Backgrounds are more common in the set of partially contained events, as is expected, and are most prevalent in backwards facing directions because of the low signal statistics in the region. For fully contained events the largest source of background is neutral current interactions, while for partially contained events it is cosmic rays. In the breakdown of signal interaction channels, separation between QE, RES, and DIS is apparent over each of the three kinematic variables. QE interactions are most common at low energies, but when the do occur at high E_ν they are more forward facing and at higher P_μ than other interaction types, owing to the lower magnitude of energy imparted into the hadronic system. By comparison, DIS events require high energies and leave little energy for the muon exiting the interaction.

The unfolded analysis binning is largely (but not completely) grid-like, constructed by taking a completely grid-based binning and merging certain bins with low statistics. In this manner, the detector resolutions, discussed more in Sec. 7, are considered in the construction of the grid slices, and the available statistics as distributed over the phase space is considered through the merging of bins. There are a total of 138 analysis bins, constructed from the familiar slice definitions of $E_\nu \in \{0.2, 0.57, 1.05, 1.57, 4.0\}$ GeV, $cos(\theta_\mu) \in \{-1, -0.5, 0, 0.27, 0.45, 0.62, 0.76, 0.86, 0.94, 1\}$, as well as nine muon momentum slices defined by the edges $P_\mu \in \{0, 0.18, 0.3, 0.45, 0.61, 0.77, 0.97, 1.28, 1.66, 2.5\}$ GeV/c.

The selection efficiency over the 3D distribution is shown in Fig. 64, showing that a non-

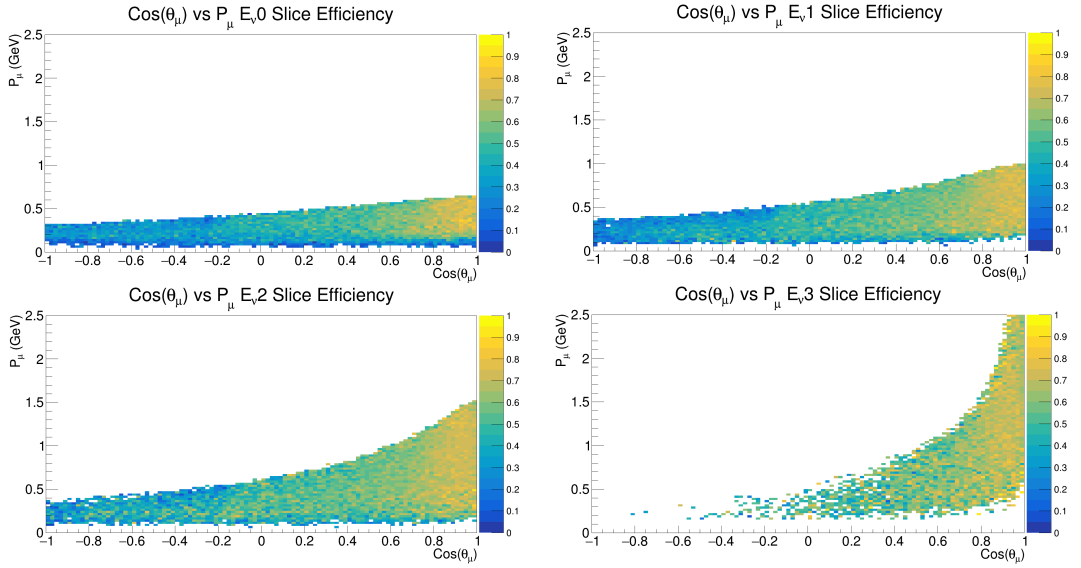

Figure 64: Selection efficiency for simulated events over $\{P_\mu, cos(\theta_\mu)\}$ within each E_ν slice. E_ν energy bins are $\{[0.2, 0.705], [0.705, 1.05], [1.05, 1.57], [1.57, 4.0]\}$ GeV.

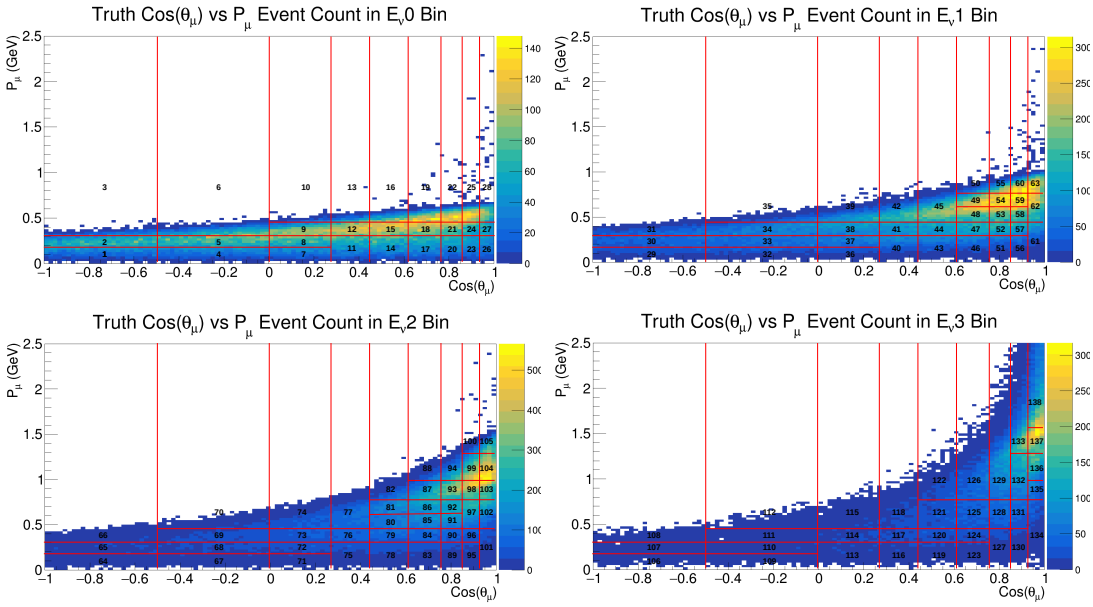

Figure 65: Distribution of simulated events over truth $\{P_\mu, cos(\theta_\mu)\}$ for each of the four E_ν slices, determined by the bin edges [0.2, 0.705, 1.05, 1.57, 4.0] GeV. The 138 analysis bins are shown in red.

zero reconstruction efficiency is maintained over the complete phase space of reconstructed events. Regions with very few simulated events are not drawn because of the inaccurate efficiency estimation in these spaces. Since (almost) no interactions are expected to exist at

those phase space regions, the efficiency there is not relevant. The distribution of simulated events over the 3D distribution is shown in Fig. 65. The red lines indicate the analysis bins, constructing by appropriately merging bins from the grid structure.

7 Resolutions of Reconstructed Quantities

The Wire-Cell reconstruction of the highly detailed LArTPC detector measurements has enabled the high-performance event selection of the previous selection, and also allows for precision measurement of reconstructed variables used in cross section measurements. However, difficulties with the MicroBooNE LArTPC such as non-functional wires, as well as imperfections in the reconstruction, leave room for improvement on future LArTPC experiments and analyses. The following sections detail both the success and limitations of the Wire-Cell reconstruction, ultimately in terms of the resolution of kinematic variables suitable for use in cross section measurements. Sec. 7.1 gives details on the foundation of the Wire-Cell reconstruction, the capability to accurately reconstruct charge in 3D through tomographic imaging. Sec. 7.2 quantifies the performance of the neutrino vertex identification algorithms at the heart of the particle flow tree construction. Finally, Sec. 7.3 shows the resolution of each of the kinematic variables used in producing unfolded cross section measurements.

7.1 Performance of Charge Reconstruction

The success of charge reconstruction through 3D imaging 5.2 has already been discussed. However, the precise accuracy in charge reconstruction has not been thoroughly quantified. Errors in charge reconstruction can contribute to difficulties in the later trajectory fitting and pattern recognition tools, impacting the overall resolution of kinematic quantities such as particle energies.

In this section, the charge reconstruction performance is evaluated through the metrics of purity and completeness to give context to further statements on detector resolutions. Charge reconstruction purity is defined as the number of reconstructed hits overlapping with truth hits divided by the total number of reconstructed hits. Charge reconstruction completeness is defined as the number of reconstructed hits overlapping with truth hits

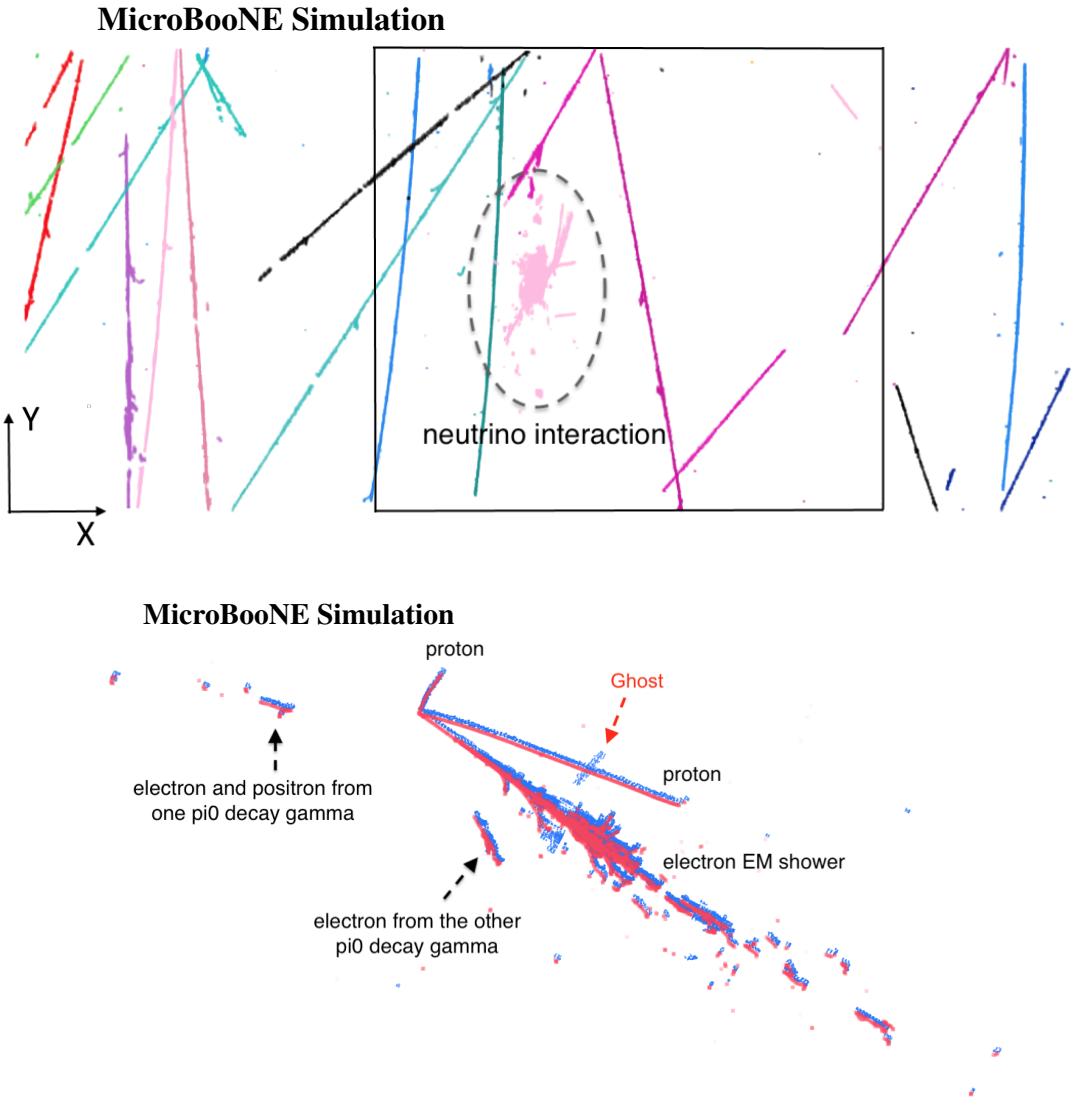

Figure 66: Event display of a $1e2p1\pi^0$ ν_eCC interaction simulated using overlay. Top: projected xy view of the TPC readout, with each cluster labeled by color. The black box represents the drift window corresponding to the beam spill timing. Bottom: Reconstructed (red) and truth (blue) charge distributions are shown. The reconstructed charge is artificially offset by ~ 1 cm for visual clarity. Source: [106]

divided by the total number of truth hits. These terms can be expressed through the set of reconstructed hits, R and the set of true hits T:

$$\begin{aligned} \text{Charge Reconstruction Purity} &= \frac{|R \cap T|}{|R|} \\ \text{Charge Reconstruction Completeness} &= \frac{|R \cap T|}{|T|} \end{aligned} \tag{7.1}$$

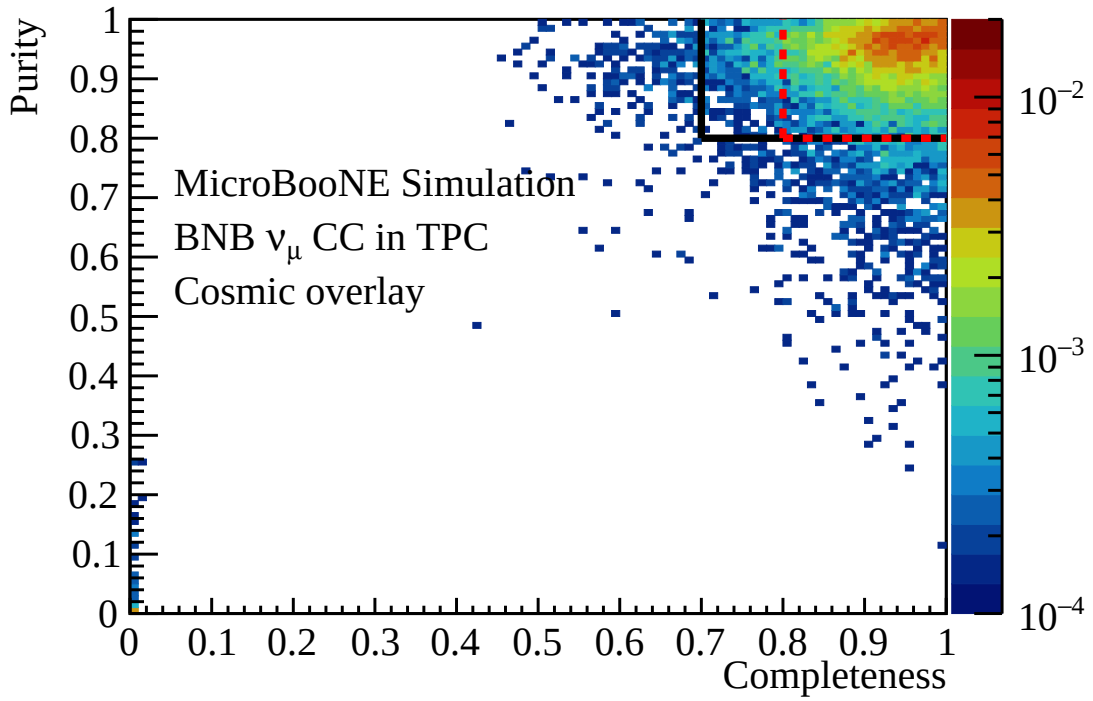

Figure 67: Distribution of purity and completeness scores on overlay ν_μCC events after imaging and charge-light matching. The black (red) box bounds a region of well reconstructed events with > 80% purity and > 70% (> 80%) completeness, totaling 80.2% (73.0%) of events. Source: [106]

Fig. 66 shows an overlay event display with a simulated neutrino interaction and multiple cosmic rays, and compares the distribution of reconstructed charge to the truth distribution. Fig. 67 shows the distribution of events over their purity and completeness scores. The black (red) box highlights well reconstructed events with > 80% purity and > 70% (> 80%) completeness, totaling 80.2% (73.0%) of events. It is worth noting that errors in purity can be corrected in stages of reconstruction after imaging (when these metrics are computed) by ignoring or rejecting areas of faulty charge reconstruction (such as a trajectory fit ignoring ghosts in a region with non-functional wires). However, errors in completeness can not be so easily recovered, as missing charge cannot be reconstructed by later steps.

7.2 Vertex Identification Resolution

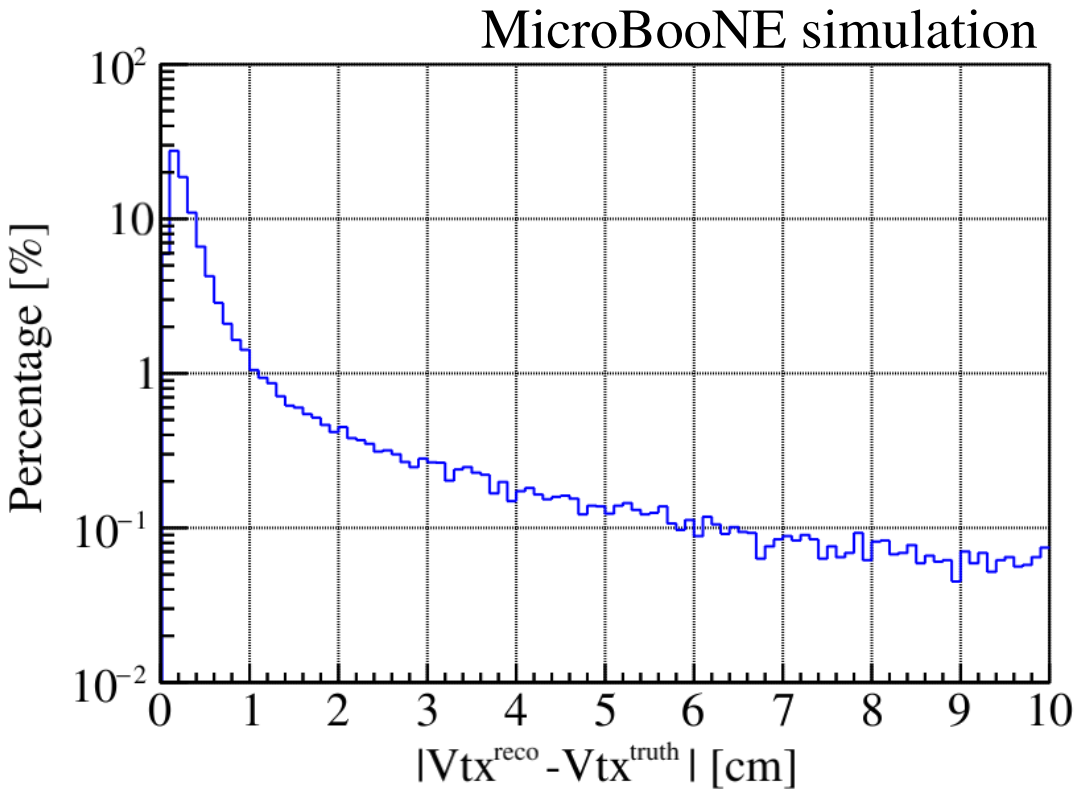

Figure 68: Neutrino vertex reconstruction of GND-selected overlay events. The distribution of events is shown as a function of the reconstructed vertex displacement magnitude. Statistical uncertainties are shown. Source: [125]

Neutrino vertex reconstruction is highly influential in both event selection and overall reconstruction performance. There are essentially two types of vertexing errors: categorical errors where the incorrect particle interaction is labeled as the neutrino interaction vertex, and precision errors where the interaction is identified correctly, but there are small positional errors in the identified location of the vertex. Categorical errors are indicative of a poorly reconstructed particle flow tree, while precision errors can lead to poor trajectory fitting, dE/dx reconstruction, and particle identification.

The distribution of GND-selected overlay events as a function of the reconstructed vertex displacement magnitude is shown in Fig. 68. To help separate categorical and precision error

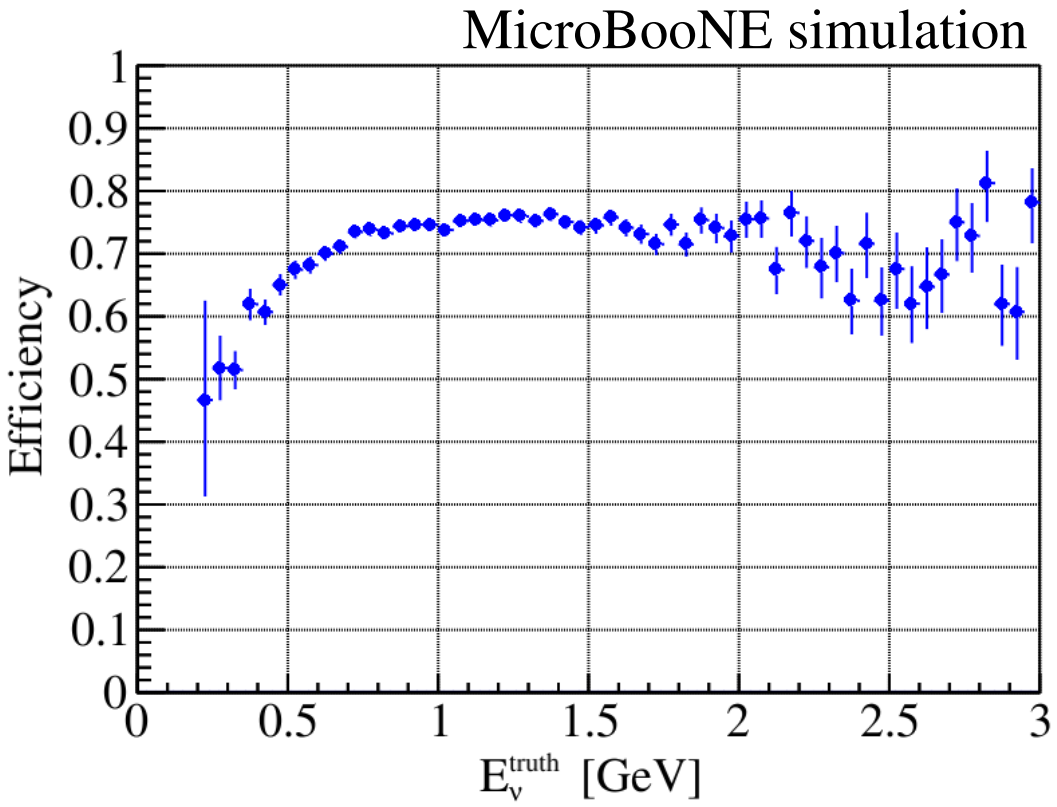

Figure 69: Neutrino vertex reconstruction of GND-selected overlay events. The reconstruction efficiency (percentage of events with a reconstructed ν vertex displacement magnitude < 1 cm) is shown as a function of the true neutrino energy. Statistical uncertainties are shown. Source: [125]

cases, an efficiency metric is defined, labeling vertices as successfully reconstructed if the displacement magnitude is within 1 cm. The reconstruction efficiency is shown as a function of neutrino energy in Fig. 69. The decrease in efficiency below 500 MeV is likely a result of the reduced effectiveness of heuristics such as “the neutrino interaction vertex is located upstream of the TPC activity, resulting from the imparted momentum from the neutrino." As discussed in Sec. 5.5, these heuristics are incorporated into the particle flow tree construction, but not strictly relied on.

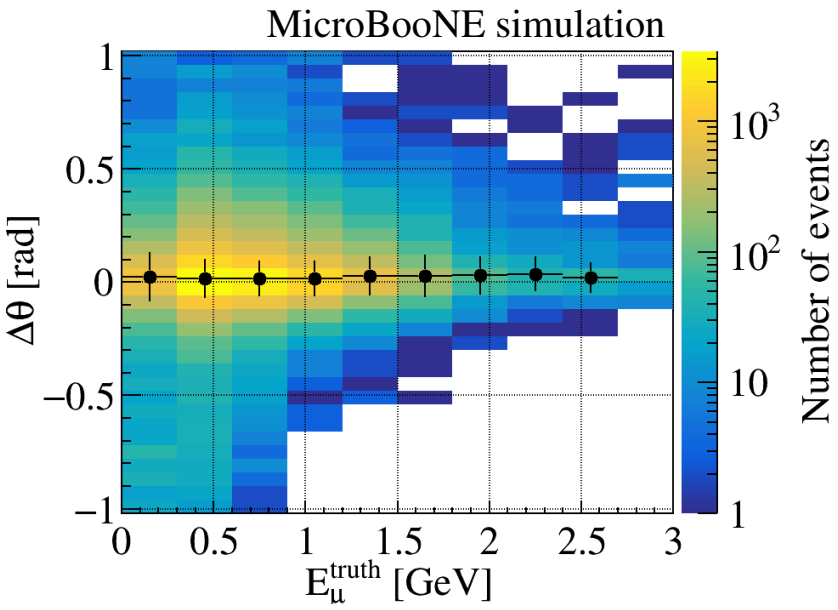

Figure 70: Difference between reconstructed and truth leading-muon-track scattering angle for selected overlay ν_μCC events, as a function of truth muon energy. Gaussian distributions are fit to each E_μ bin, with the fit parameter values overlaid. Since the polar angle cannot be negative, it has a slight positive bias overall. Source: [125]

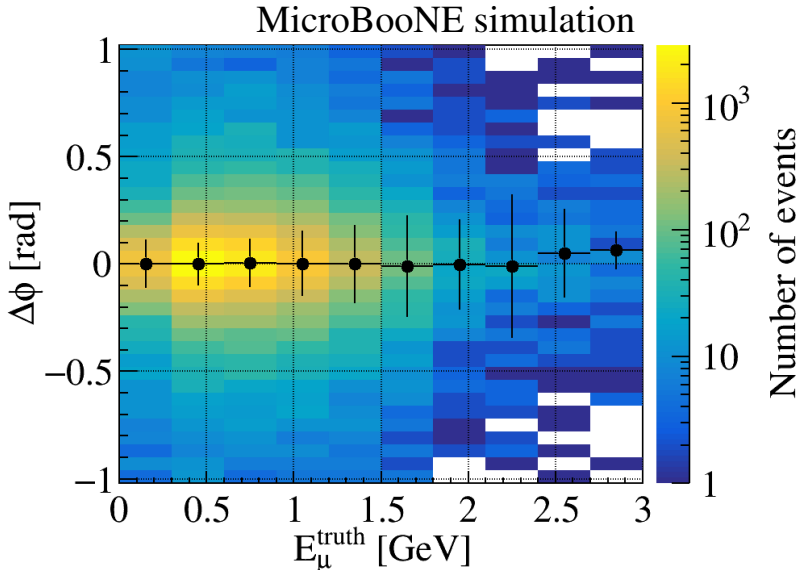

Figure 71: Difference between reconstructed and truth leading-muon-track azimuthal angle for selected overlay ν_μCC events, as a function of truth muon energy. Gaussian distributions are fit to each E_μ bin, with the fit parameter values overlaid. Source: [125]

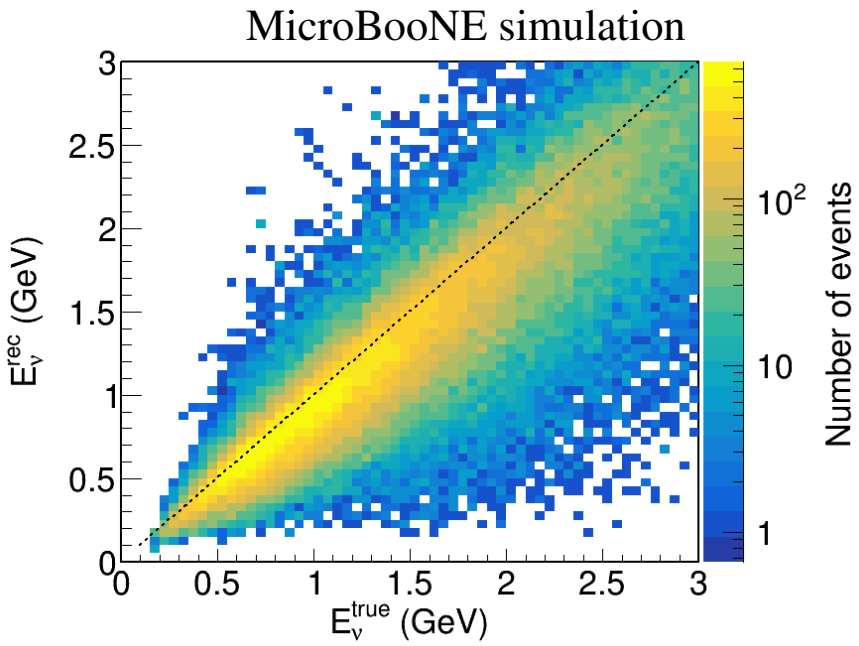

Figure 72: Reconstructed vs truth neutrino energy for selected overlay ν_μCC events. Source: [125]

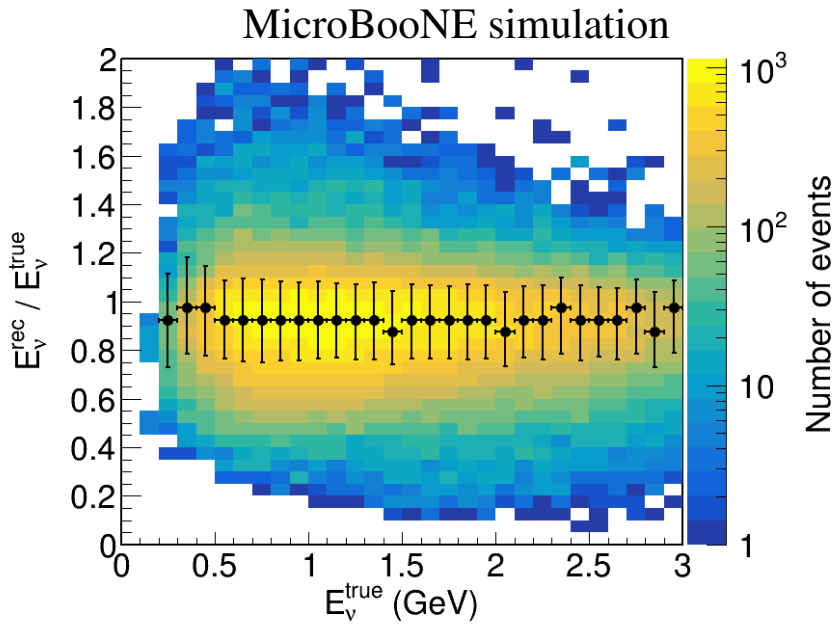

Figure 73: Difference between reconstructed and truth neutrino energy for selected overlay ν_μCC events, as a function of truth neutrino energy. Gaussian distributions are fit to each E_ν bin, with the fit parameter values overlaid. Source: [125]

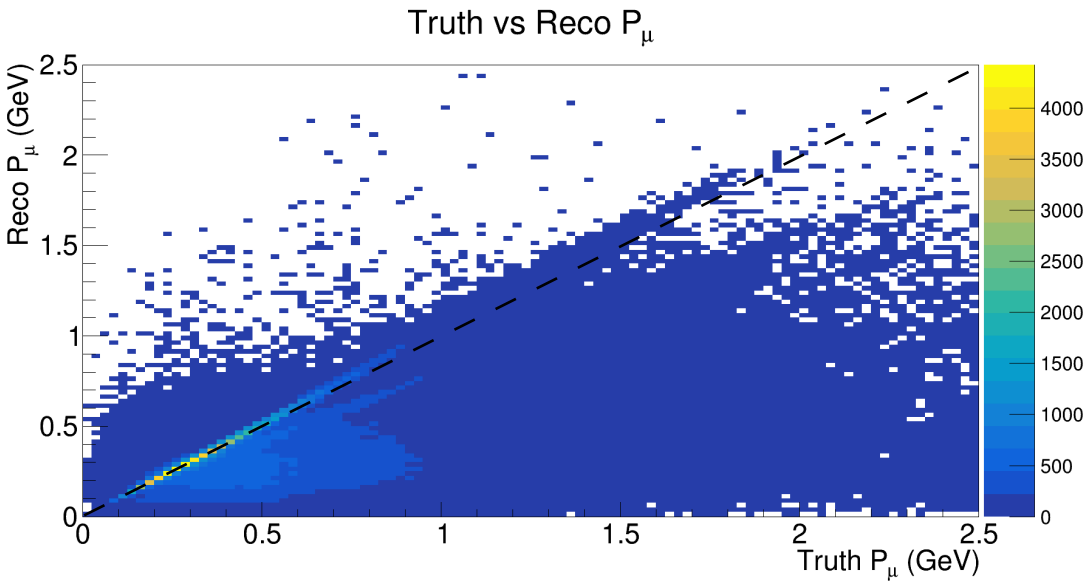

Figure 74: Reconstructed vs truth muon momentum for selected overlay ν_μCC events.

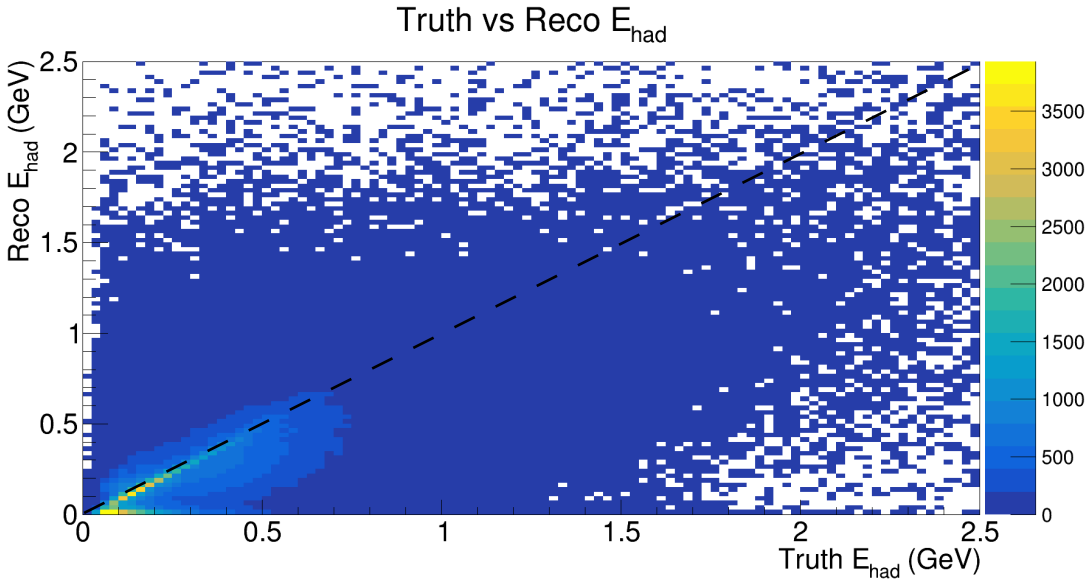

Figure 75: Reconstructed vs truth visible hadronic energy for selected overlay ν_μCC events.

7.3 Resolutions of Kinematic Variables

For each kinematic variable, the resolution is computed by comparing the reconstructed value to the truth value using overlay events. The distribution of events is examined over bins of the corresponding truth variable, and a Gaussian distribution is fit to determine the reconstructed resolution [125, 137]. Fig. 70 and Fig. 71 show the measured angular resolution

for the muon scattering angle θ_μ and the muon azimuthal angle ϕ_μ, respectively. At forward scattering angles, the resolution is $\sim 5^\circ$, performing worse at backward angles. Fig. 72 and Fig. 73 show the measured neutrino energy resolution, which is seen to be $\sim 20\%$ across the energy range. Fig. 74 shows the measured muon momentum resolution, and Fig. 75 shows the measured visible hadronic energy resolution, both of which are approximately $10-15\%$. The sub-dominant off-diagonal band of reconstructed muon momentum is the result of imperfect energy reconstruction using the summation of dE/dx method discussed in Sec. 5.5.7. In future analyses this method could be improved, and augmented with other approaches such as multiple Coulomb scattering.

8 Estimation of Systematic Uncertainties

The analysis incorporates statistical uncertainties from the number of events as well as systematic uncertainties from a wide range of estimated parameters into the overall measurement. This is done through the covariance matrix formalism [141], which describes the correlated uncertainties between each of the reconstructed bins (and after unfolding, between analysis bins). Because of the bilinearity of covariance, contributions from each systematic effect can be computed separately and summed to form the overall covariance matrix. Sec. 8.1 describes the estimation of statistical uncertainties on the sample of simulated overlay events used. Sec. 8.2 describes the estimation of uncertainties from the flux prediction, POT measurement, neutrino interactions outside the cryostat, and number of target nuclei in the fiducial volume. Sec. 8.3 describes the estimation of uncertainties from the modeling of ν-Ar interactions in the Genie event generator and hadronic interactions in the argon medium using GEANT 4. Sec. 8.4 describes the estimation of systematic uncertainties from the modeling of detector effects. Finally, the overall covariance matrix is discussed in Sec. 8.5.

8.1 Estimation of Monte-Carlo Statistical Uncertainties

It is tempting to use the frequentist approach to estimate the statistical uncertainties in the Monte-Carlo (MC) overlay event sample. Under this method, a measurement of n events in a particular bin would have an estimated uncertainty of $\sqrt{n}$. However, this approach struggles in the particular case of bins without any measured events, assigning an uncertainty of 0. Therefore, a Bayesian approach is used to predict MC statistical uncertainties included in the overall covariance matrix. The actual implementation accounts for the non-uniform weighting of simulated events, following the method described in reference [142], however the details of the simpler uniform weights case is discussed below.

The number of events measured in a given bin is assumed to be Poisson distributed,

described by a mean event rate μ. The goal, then, is to estimate μ and its uncertainty given a measurement of N events over a particular time interval. To achieve this, an uninformed prior is selected, describing the distribution of μ through a step function with value 0 for $\mu < 0$ and 1 for $\mu \geq 0$. This leads to a posterior distribution approximated as:

$$P(\mu|N) \approx exp\left(N - \mu + Nln\left(\frac{\mu}{N}\right)\right) \tag{8.1}$$

Therefore, in the case of $N = 0$, the posterior distribution becomes $P(\mu|0) = e^{-\mu}$, allowing an RMS uncertainty of $\sqrt{2}$ to be computed for a taken central value of $\mu = N = 0$. In general, uncertainties for all measured values can be computed in this manner.

8.2 Estimation of Flux Uncertainties

The flux prediction uncertainties follow the flux prediction computed for MiniBooNE [89], and primarily consider the production rate of hadrons through $p - Be$ interactions in the target, determined through GEANT4 simulation [143]. The production rate of pions is re-weighted based on hadron production data measured in HARP [144]. There are also other effects considered, such as the mis-modeling and mis-calibration of the horn current, as well as the total, QE, and inelastic cross sections of pions and nucleons on the beryllium target and aluminum horn. The flux covariance matrix is constructed by using a multisim technique, where each parameter is varied simultaneously following a Gaussian distribution given by a predicted central value and uncertainty. In total, $N = 1000$ universes are constructed this way, and the covariance Cov_{ij} between different bins x_i and x_j is computed using the formula:

$$Cov_{ij} = \frac{1}{N}\sum_{n=1}^{N}(x_{i,n} - \mu_i)((x_{j,n} - \mu_j)) \tag{8.2}$$

where $\vec{\mu}$ describes the central value prediction for each bin. An additional 2% overall normalization uncertainty is applied to account for the uncertainty in the number of protons on

Table 3: Contribution of different sources to overall flux uncertainty, given as a percentage of the central value prediction. Source [90].

Systematic	ν_μ (%)	$\overline{\nu}_\mu$ (%)	ν_e (%)	$\overline{\nu}_e$ (%)
POT	2.0	2.0	2.0	2.0
π^+	11.7	1.0	10.7	0.03
π^-	0.0	11.6	0.0	3.0
K^+	0.2	0.1	2.0	0.1
K^-	0.0	0.4	0.0	3.0
K^0_L	0.0	0.3	2.3	21.4
Other	3.9	6.6	3.2	5.3
Total	12.5	13.5	11.7	22.6

target. This includes both the uncertainty in the measured beam intensity as measured in the toroids upstream of the target, as well the uncertainty in the number of protons that actually interact with the target. The overall uncertainty contribution from each source (POT count, hadronic production, other considerations) is shown in Table 3, and the total flux correlation matrix is shown in Fig. 76.

There is also a ~ 1% uncertainty added to the overall covariance matrix because of the estimation of the number of target nuclei in the fiducial volume, resulting from the following considerations. The fiducial volume has dimensions 2.50 m × 2.26 m × 10.31 m, totaling 58.2511 m^3, so a ±1 cm uncertainty on each detector boundary contributes an overall volume uncertainty of ~ 0.35%. There is also a ~ 0.9% variation in the density of argon between the temperatures of 87 K and 89 K. Finally, the purity of argon remains above 99.6%, allowing for an overall target nuclei estimation of $1.21 \times 10^{31} \pm 1\%$. Finally, in addition to the modeling of signal neutrino interactions inside the fiducial volume, background neutrino interactions outside the cryostat, called dirt interactions, must be considered. The most difficult part of this estimation is the modeling of interactions with outside materials, therefore a conservative 50% bin-to-bin (uncorrelated) uncertainty is applied to the predicted rate of dirt events. However, since dirt is a negligible contribution to the overall background, this has a minimal

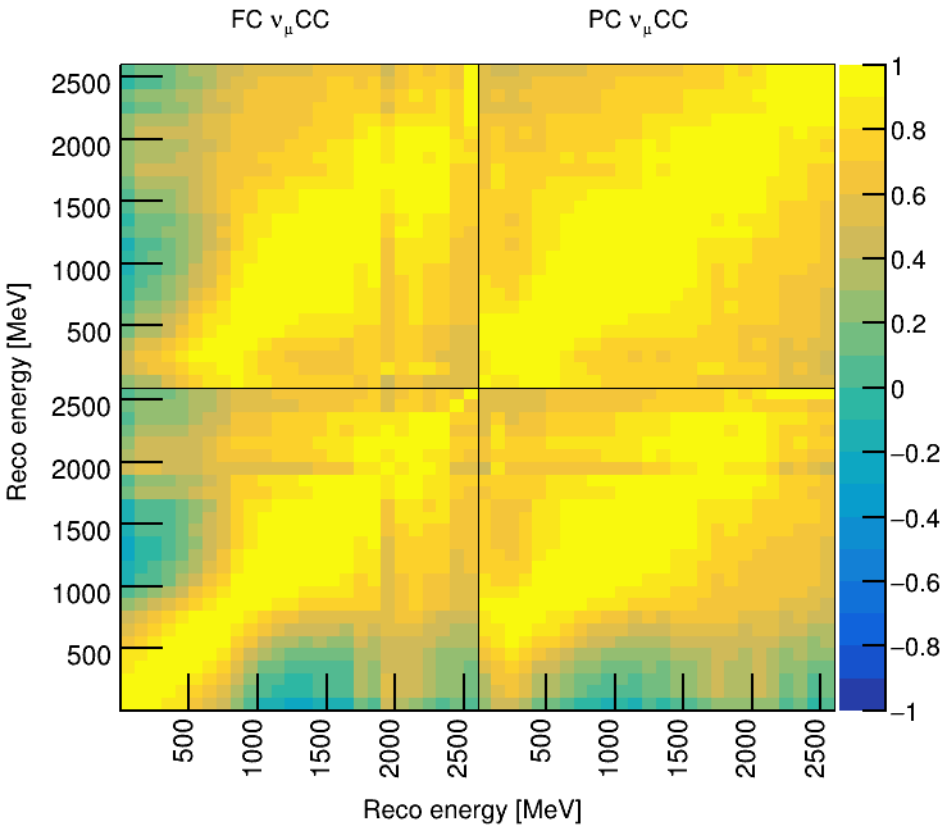

Figure 76: Correlation matrix of flux systematics across reconstructed neutrino energy, with FC and PC selections separated.

impact.

8.3 Estimation of Cross Section Uncertainties

Within the MicroBooNE experiment, neutrino interactions are simulated using the Genie event generator [145] for the central value MC prediction. Previous Genie models have been validated against bubble chamber data in the past, as well as MiniBooNE CC0π data more recently. However, MicroBooNE is the first experiment to adopt Genie v3 for its central value MC prediction, and initially found a data/MC discrepancy in the CC0π channel. As a result, MicroBooNE decided to re-tune the CCQE and CCMEC predictions to better describe the data seen in MicroBooNE [146]. To avoid bias from tuning on MicroBooNE data, T2K CC0π data

was used [147] to create a “MicroBooNE Tune" version of Genie 3.0.6 G18_10a_02_11a [148]. This version features an initial state described by the Local Fermi Gas model, and final state interactions described through an effective hA2018 intra-nuclear cascade model. For more details on the nuclear effects considered in a neutrino event generator, see the discussion in Sec. 2.3.

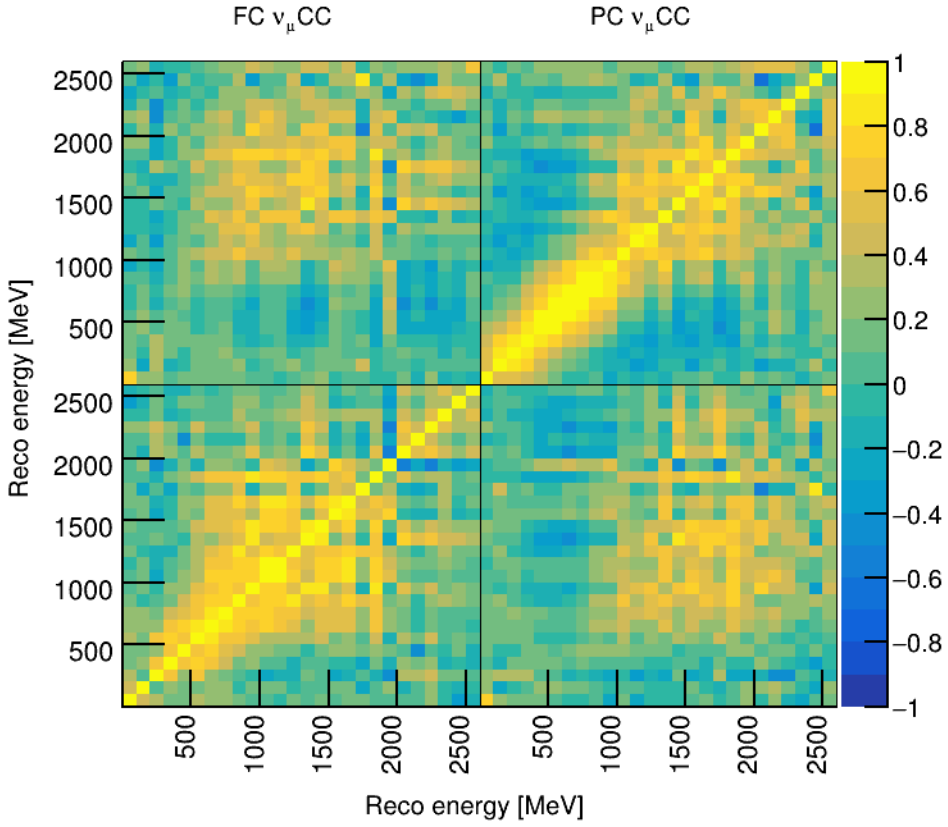

Figure 77: Correlation matrix of Genie systematics across reconstructed neutrino energy, with FC and PC selections separated.

Although cross section uncertainties are naturally suppressed in the production of cross section measurements, as the unfolded signal distribution does not include cross section uncertainties, it is still very important to accurately model the data within listed uncertainties, so as to not introduce significant bias in the unfolded result. The sufficiency of the Genie model and its uncertainties to describe the distribution in data seen in MicroBooNE will be discussed more in Sec. 9, but for now the estimation of Genie uncertainties will be briefly

discussed. Within Genie, there are a large number of parameters available for re-weighting that are used in the simulation of neutrino interactions, with a full description available in the Genie users manual [149]. To estimate the overall uncertainty contribution from the choice of each parameter value, they are varied simultaneously in a multisim approach, as described in Sec. 8.2. A total of 600 universes are generated and the resulting correlation matrix across neutrino energy is shown in Fig. 77.

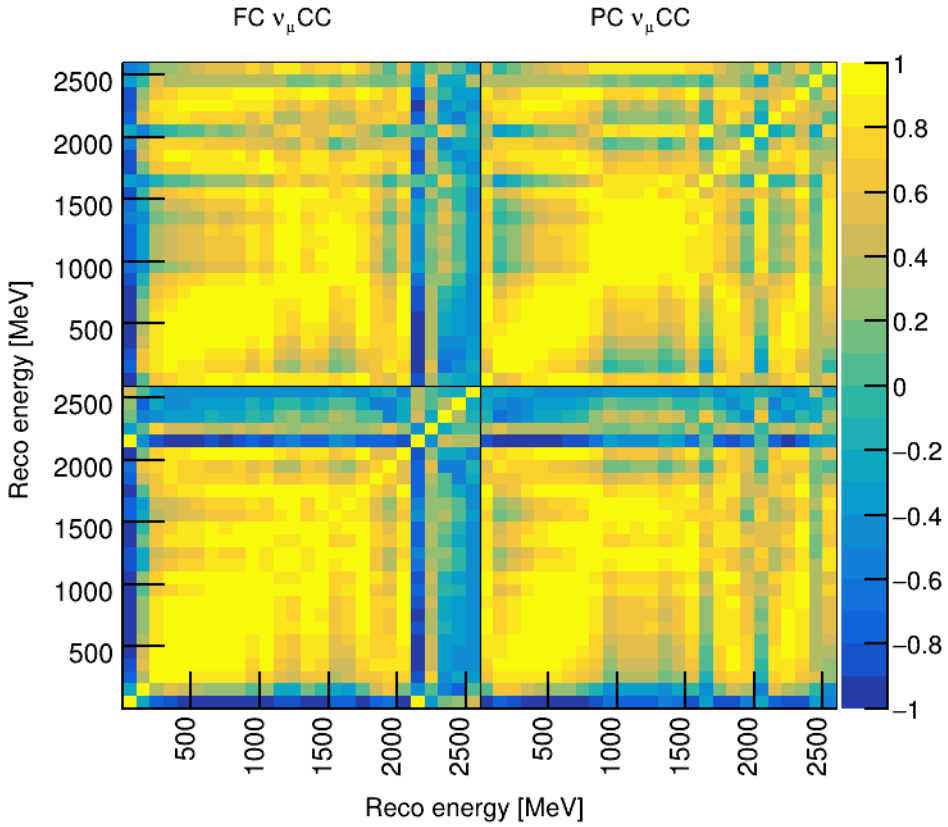

Figure 78: Correlation matrix of GEANT4 systematics across reconstructed neutrino energy, with FC and PC selections separated.

The interaction between hadronic final states and external argon nuclei throughout the detector volume is modeled using GEANT4 [143]. In addition to ionizing electrons, interactions with atomic nuclei can occur that can change an ionizing particle's direction, or form new particles. The chance of interaction is described by the hadron-nucleus cross

section, and for a given particle energy and travel distance, survival and (complementary) re-interaction probabilities are computed. These are used to assign weights to interacting and non-interacting hadrons on a per-event basis. A conservative 30% uncertainty is assigned to the interaction cross section, contributing an overall 0.612% uncertainty on the neutrino interaction cross section as determined through multisim using 10 universes. The correlation matrix across neutrino energy is shown in Fig. 78

8.4 Estimation of Detector Systematic Uncertainties

The detector response simulates many non-idealizations of the MicroBooNE detector that can impact the overall reconstruction and selection of events. The different effects considered can be grouped into three categories: light yield and propagation throughout the detector, recombination and space charge effects on the drift of ionization charge, and variations in the reconstructed TPC waveform. For each effect, a 1σ uncertainty is defined so that the impact on the overall event selection can be determined through a unisim procedure [150].

Light yield considers the reduction in PEs observed at the PMT array resulting from attenuation through the detector as well as degradation of the PMT array. Over the course of runs 1-3 of data taking used in this analysis there was a $\sim 25\%$ reduction in the measured PE resulting from PMT degradation, therefore this difference is taken as the 1σ uncertainty value on the simulated light yield. The 1σ uncertainty on the recombination effect is estimated by taking a different set of parameters on the modified box recombination model and computing the difference in predicted recombination. The space charge effect is mapped by studying throughgoing muon tracks, assumed to be largely straight, and noticing the common deviations in path at each voxelized detector location. The 1σ uncertainty on this map is taken as the residual deflection from straight-line tracks after applying the space charge map to tracks produced by a laser shined through the detector volume. The TPC waveform 1σ uncertainty is computed as the ratio of measured to simulated values of reconstructed charge and signal

width after deconvolution and digitization. By using the digitized values, the estimation is computationally efficient and also captures the entire range of non-idealizations in the simulated TPC waveform that are not directly simulated [151].

A sample of overlay events is produced using parameter central-values, and additional samples are created after varying each parameter by the previously described 1σ values one at a time (unisim). Because of random processes in the event processing and reconstruction software, there may be uncorrelated uncertainties for each simulated event, in addition to the correlated uncertainties between bins. It is important to understand the divide between correlated and uncorrelated uncertainties both for an accurate treatment of the detector uncertainties, and particularly in the case of the conditional constraint procedure used for model validation in Sec. 9. This procedure will be discussed in more detail in the following section, but essentially it relies on a Bayesian procedure to update the model prediction over a particular distribution of interest by using the data measurement over another distribution of a more trusted reconstructed variable. Correlations between predictions allows the constraint procedure to limit the posterior prediction over the distribution of interest, but it is important to correctly identify the uncorrelated part of the uncertainty so that the updated prediction is correctly computed. Over a sufficiently large sample, these variations may be fully captured; however, computational constraints and the use of overlay events limits the available sample size. Since each simulated neutrino interaction is paired with an EXTBNB readout of cosmic ray activity, the maximum number of overlay events is determined by the amount of EXTBNB data available.

8.4.1 Bootstrapping

To compensate for the limited sample size available, a bootstrapping procedure is used to re-sample the reconstruction and selection of events many times. This allows for the correct determination of correlated variations between bins, as well as uncorrelated variations

resulting from random processes in the reconstruction, as opposed to statistical fluctuations. The set of all overlay events is sampled once using the central values for each detector response parameter, and again with the 1σ deviation for each detector response parameter (applied one at a time), in what is described as one universe of simulation. The difference in selected event distribution between the CV and 1σ sample for each parameter is taken as a difference vector $\vec{V}_{D,ij}$ representing the 1σ uncertainty in event selection distribution associated with the uncertainty in the i'th detector response parameter for the j'th universe. This process is repeated $N_1 = 1000$ times to capture the range of possible variations in $\vec{V}_{D,i}$ resulting from the random processes in the event processing and reconstruction. The average 1σ deviation $\vec{V}_{D,i}^{nominal}$ is computed by averaging the difference vectors across all simulated universes:

$$\vec{V}_{D,i}^{nominal} = \frac{1}{N_1} \sum_{j=1}^{N_1} \vec{V}_{D,ij} \tag{8.3}$$

Additionally, the uncertainty on $\vec{V}_{D,i}^{nominal}$ is computed as the covariance matrix $M_{R,i}$ from the difference vectors for each universe, $\vec{V}_{D,ij}$, following the familiar covariance matrix construction described in Eq. 8.2. Since $\vec{V}_{D,i}^{nominal}$ describes the uncertainty on the overall event selection from the i'th detector response parameter, $M_{R,i}$ describes an uncertainty on an uncertainty.

The overall covariance matrix is constructed from the both $\vec{V}_{D,i}^{nominal}$ and $M_{R,i}$ in a way that fully captures the detector response uncertainty. Again $N_2 = 1000$ universes are used, although these are separate from the universes discussed earlier. For each universe, an overall difference vector $\vec{V}_{D,i}^{final}$ is constructed from $\vec{V}_{D,i}^{nominal}$ as well as a deviation $\delta\vec{V}_{D,i}$ representing the uncertainty in $\vec{V}_{D,i}^{nominal}$. The deviations are constructed from the n eigenvectors $\{\vec{e}_{1,i}, ..., \vec{e}_{n,i}\}$ and eigenvalues $\{\lambda_{1,i}, ..., \lambda_{n,i}\}$ as well as n Gaussian distributed scalars $r_{k,i} \in N(0,1)$ following:

$$\delta\vec{V}_{D,i} = \sum_{k=1}^{n} r_{k,i}\sqrt{\lambda}_{k,i}\vec{e}_{k,i} \tag{8.4}$$

There is also an overall Gaussian distributed scalar $r_{0,i} \in N(0,1)$ that scales the entire difference vector to construct a sampled event distribution for the universe:

$$\vec{V}_{D,i}^{final} = r_{0,i}\left(\vec{V}_{D,i}^{nominal} + \delta\vec{V}_{D,i}\right) \tag{8.5}$$

After repeating the process of generating random numbers $r_{k,i} \in N(0,1) \forall k \in [0,n]$ for each of the N_2 universes, the overall covariance matrix $M_{D,i}$ for the i'th detector variation can be constructed from the $\vec{V}_{D,i}^{final}$ of each universe following Eqn. 8.2 for constructing a covariance matrix. Finally, the complete detector covariance matrix M_D is simply the sum of covariance matrices for each detector response effect considered.

8.4.2 Gaussian Processes Regression Smoothing

The procedure above is sufficient for analyses with a small number of bins, where the number of events per bin under the bootstrapping procedure is not too small. However, in the case of a triple-differential analysis there are enough bins, covering rare regions of the kinematic phase space, that statistical fluctuations become extreme in the computation of $M_{R,i}$. This leads to over-estimation of the uncertainty on $\vec{V}_{D,i}^{nominal}$ and ultimately on the total covariance matrix M_D. Therefore, to reduce the over-estimation of statistical fluctuations, a smoothing technique is applied within the bootstrapping procedure to the computation of $\vec{V}_{D,i}^{nominal}$ and $M_{R,i}$ before $M_{D,i}$ is computed. Smoothing describes a class of algorithms that assert an intuition of smoothness on a function, de-preferencing extreme variations in the construction of a function to describe the data. In particular, Gaussian Processes Regression (GPR) [152, 153, 154] smoothing is employed in this analysis, which has been used in physics broadly [155], as well as in high energy physics in particular [156] [157] [158]. GPR assumes measurement bins to be jointly Gaussian distributed with a particular covariance, and uses Bayesian statistics to form a prediction based on the measured data. The GPR posterior distribution takes the form of non-parametric curves that are able to describe the observed

data. The mean posterior prediction from GPR smoothing is used as the as the smoothed central value prediction, $\hat{V}_{D,i}^{nominal}$, with a smoothing error estimate given by the covariance of the posterior, $\hat{M}_{R,i}$. In this manner, GPR smoothing is able to take as input the simulated estimates for the nominal difference vector $\vec{V}_{D,i}^{nominal}$ and its uncertainty $M_{R,i}$ and reduce the statistical fluctuations to form a posterior prediction that better describes the overall detector covariance from the stated detector response parameter uncertainties. Smoothing is performed on each detector response parameter separately; going forward the index over parameters will be ignored.

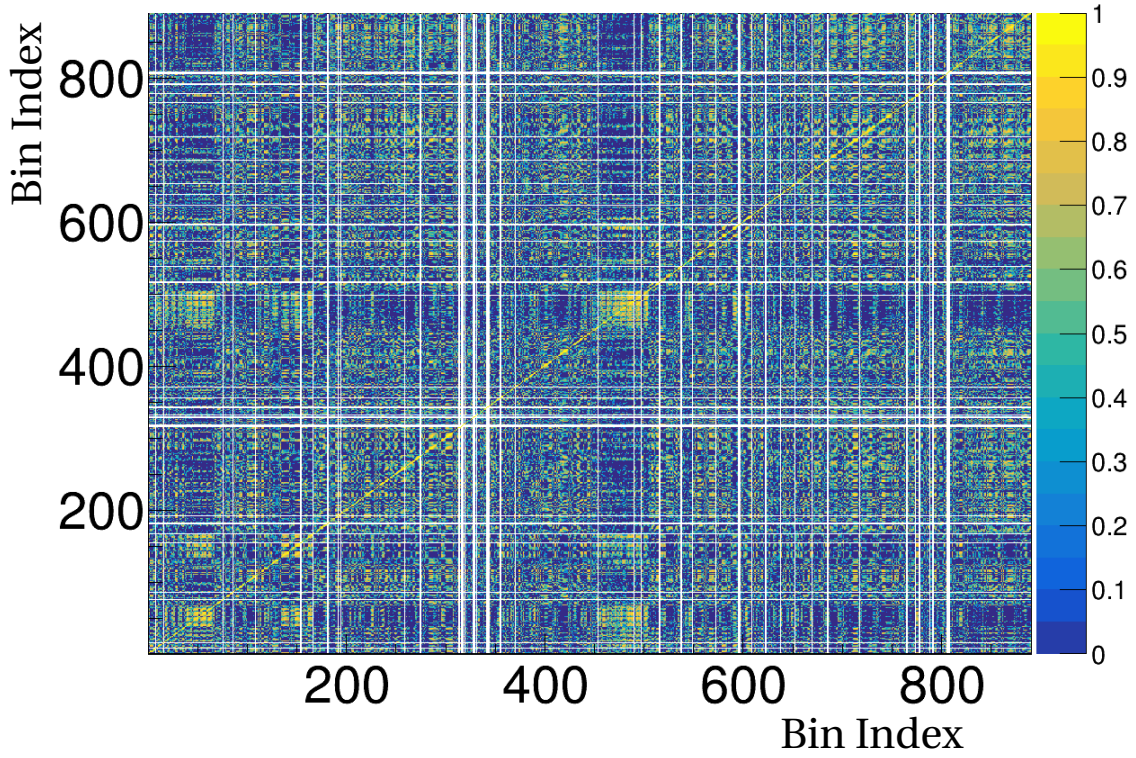

Figure 79: Correlation matrix of detector response uncertainties across reconstructed bins after removing bins with zero predicted and measured value. Bins are arranged in the top-down hierarchy of FC/PC, E_ν slice, $cos(\theta_\mu)$ slice, and finally P_μ slice.

GPR begins with an uninformed prior $p(\vec{x}_a) = N(\vec{\mu}_a, \Sigma_{K,aa}) = N(\vec{0}_a, \Sigma_{K,aa})$ over the set of prediction points $\vec{x}_a$, where $\vec{\mu}_a$ represents the mean prediction at each point and $\Sigma_{K,aa}$ is a covariance matrix describing the correlations between prediction points based on a kernel function $K(x_i, x_j)$. Note that since the prediction points are located in the 3D phase

space $\{E_\nu, cos(\theta_\mu), P_\mu\}$, each prediction point is itself a vector. The kernel function asserts the intuition of smoothness by correlating nearby points highly, chosen to be a Radial Basis Function (RBF) in this analysis:

$$\Sigma_{K,aa,ij} = K(\vec{x}_i, \vec{x}_j) = e^{-|(\vec{x}_i - \vec{x}_j)\cdot\vec{s}|^2/2} \tag{8.6}$$

where $s_i = \frac{1}{L_i}$ for characteristic length scales L_i along each of the three kinematic dimensions. Length scales are chosen based on measured detector resolutions to be 0.1 in $cos(\theta_\mu)$, 20% in E_ν, and 20% in P_μ, where fractional length scales are achieved by taking the log of the energy and momentum coordinates for use within GPR smoothing. This prior is then updated through the use of the measurement $\vec{y}(\vec{x}_b) = \vec{V}_D^{nominal}$ over measurement points $\vec{x}_b$, although in this analysis the prediction points and measurement points are the same, so $\vec{x}_a = \vec{x}_b$. The measurement has uncertainty given by the covariance matrix $\Sigma_y = M_R$, which is added to the kernel function to form the total covariance matrix over the measurement points $\Sigma_{T,bb}$:

$$\Sigma_{T,bb} = \Sigma_y + \Sigma_K \tag{8.7}$$

Given the prior prediction and the measurement, a posterior prediction can be constructed:

$$\begin{aligned}
\hat{p}(\vec{x}_a|\vec{y}) &= N(\hat{\mu}_{a|y}, \hat{\Sigma}_{T,aa|y}) \\
\hat{\mu}_{a|y} &= \vec{\mu}_a + \Sigma_{K,ab}\left(\Sigma_{T,bb}\right)^{-1}(\vec{y} - \vec{\mu}_b) \\
\hat{\Sigma}_{T,aa|y} &= \Sigma_{K,aa} - \Sigma_{K,ab}\left(\Sigma_{T,bb}\right)^{-1}\Sigma_{K,ba}
\end{aligned} \tag{8.8}$$

The posterior prediction CV $\hat{\mu}_{a|y}$ and covariance $\hat{\Sigma}_{T,aa|y}$ are used in place of the originally simulated values $\vec{V}_D^{nominal}$ and M_R. Because of GPR smoothing, statistical fluctuations are controlled and become less impactful in M_D, reducing the overall detector response covariance from ~100% to ~50%. The detector response correlation matrix (after removing

bins with zero data and prediction) is shown in Fig. 79.

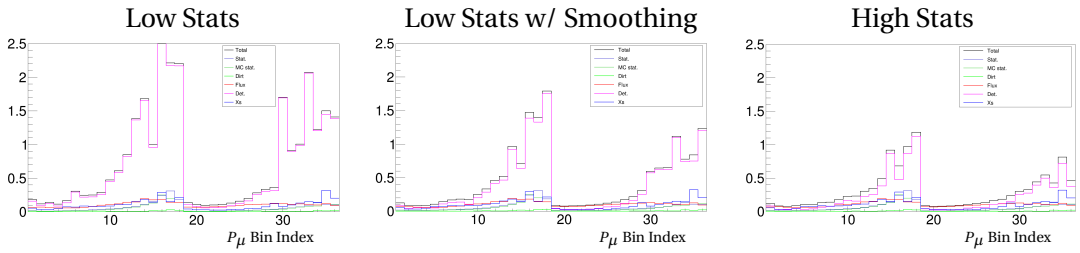

Figure 80: Diagonal covariance matrix elements for each uncertainty, stacked, over the 1D P_μ distribution, with FC and PC sections separated. Detector response uncertainties are shown in magenta, labeled as “Det." Left: low (~ 10%) stats data set used. Middle: low stats data set with GPR smoothing. Right: High stats data set used.

A test was performed over the 1D P_μ distribution to see whether GPR smoothing would reduce the detector response uncertainties beyond the point enabled by using a high-statistics data set. The test was performed by comparing the detector response uncertainties from a small data set containing only ~ 10% of events before and after GPR smoothing with the uncertainties under the full data set. By using a small data set, the conditions found in the triple-differential analysis were roughly reproduced, and statistical fluctuations caused a massive over-estimation of detector response uncertainties. After applying GPR smoothing to $\vec{V}_D^{nominal}$ and M_R, the estimated uncertainties shrunk considerably, but notably not below the values achieved through the high-stats data set, shown in Fig. 80. This outcome agreed with the premise that GPR smoothing reduces statistical fluctuations without suppressing physical uncertainties computed through the bootstrapping procedure.

8.5 Total Covariance Matrix

The overall correlation matrix over the reconstructed bins is shown in Fig. 81. The breakdown of uncertainties (diagonal elements only) between the various statistical and systematic components is shown in Fig. 82. The largest systematic is the detector response uncertainty, comprising an average of ~ 50% of the data values. This is followed by cross section and flux uncertainties, at ~ 20% and ~ 5 – 15% over the energy range, respectively. Stats and MC stats

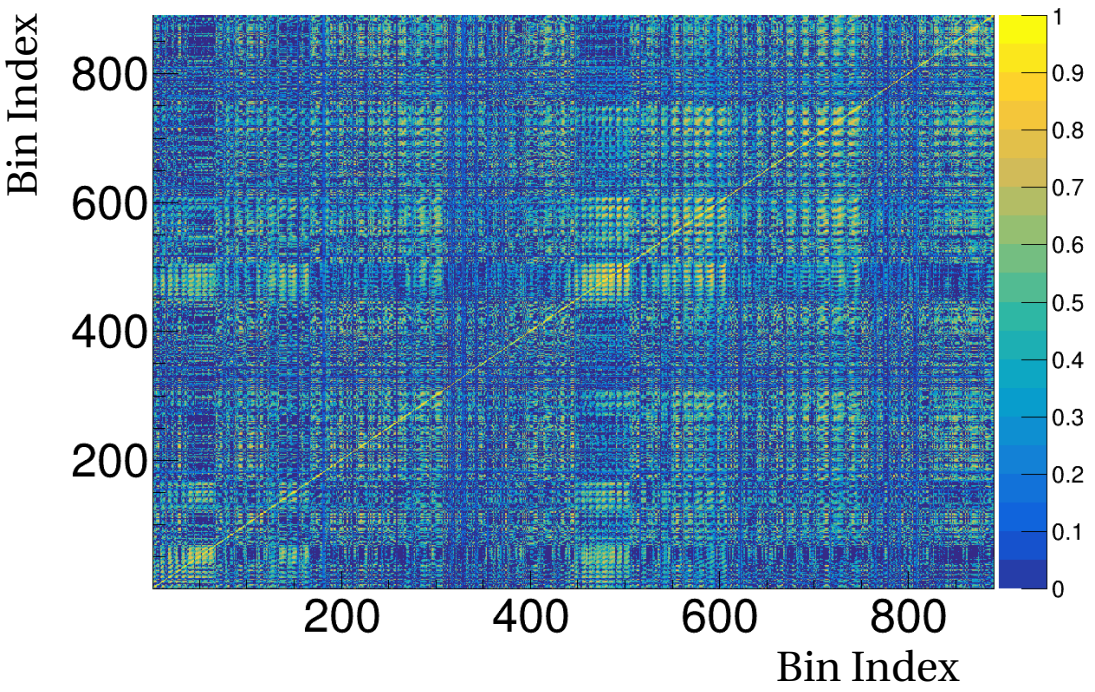

Figure 81: Correlation matrix of overall uncertainties across reconstructed bins after removing bins with zero predicted and measured value. Bins are arranged in the top-down hierarchy of FC/PC, E_ν slice, $cos(\theta_\mu)$ slice, and finally P_μ slice.

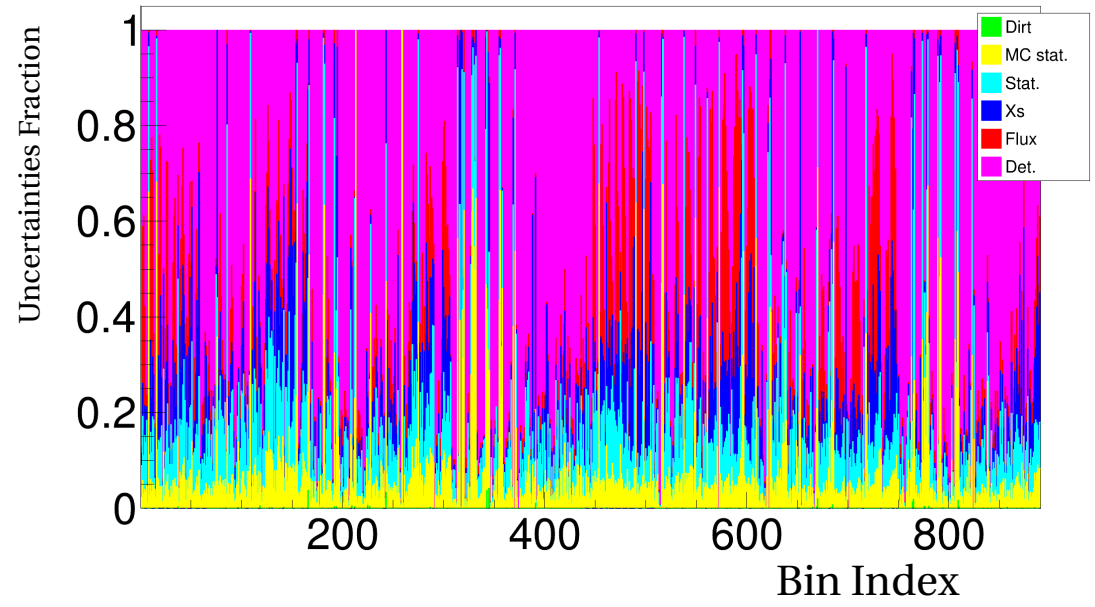

Figure 82: Breakdown of uncertainties (diagonal elements only) by source after removing bins with zero predicted and measured value. Bins are arranged in the top-down hierarchy of FC/PC, E_ν slice, $cos(\theta_\mu)$ slice, and finally P_μ slice.

each contribute ~ 20 uncertainties, and GEANT4, dirt, and number of target nuclei all have minor contributions.

9 Model Validation

Neutrino experiments rely on model predictions to convert event counts over reconstructed variables into measurements on physically relevant quantities. Event selections fail to detect some signal events, described by the selection efficiency, and incorrectly label some background activity as signal, described by the selection purity. Furthermore, even within correctly selected events, there can be biases that alter the distribution of reconstructed events, such as an under-estimation of neutrino energy after failing to reconstruct all the interaction activity. Reconstruction biases can lead to bin migration, such as to lower energy bins in the previous example, as well as bin smearing resulting from poor variable resolution, and even more complicated effects. If not properly corrected for, measured quantities will not produce an accurate statement on physical quantities of interest, such as oscillation parameters or cross sections.

There are two general ways to account for the range of biases introduced in producing an event selection: forward folding and unfolding. Forward folding leaves the measurement in terms of reconstructed variables and smears the model prediction over a truth distribution by the estimated detector response, allowing a direct comparison with the measured data. Note that here “detector response" includes all aspects of the experiment that separate a measured event count from the underlying physical parameter of interest, including the flux and cross section, unlike the more narrow “detector response" considered in Sec. 8.4. The detector response mapping includes the effects discussed earlier, such as selection efficiency, bin migration, and bin resolution, and can be described by the matrix equation:

$$\vec{M} = R \cdot \vec{S} + \vec{B} \tag{9.1}$$

where $\vec{M}$ is the measured distribution over reconstructed bins, $\vec{S}$ is the predicted truth distribution, R is the response matrix that maps from truth to reconstructed variables, and $\vec{B}$ is the

predicted distribution of selected background events. Forward folding can be advantageous when an analysis requires a coarse binning of reconstructed variables, either as a result of low available statistics or poor measurement resolution, as it is easy to map fine-grain model predictions into any reconstructed binning. Forward folding can also yield more powerful model tests [159], however it has the downside of not producing a measurement over the physically relevant truth variables.

In contrast with forward folding, unfolding accounts for selection biases by creating an inverse mapping of the detector response, and computing the measured distribution in terms of truth variables after undoing the detector response effects. This provides a result that is more physically intuitive and relevant, and can be directly compared with other model predictions or experimental measurements on the same physics. The details of unfolding are more complicated than simply computing an inverse matrix to R, and will be discussed in more detail in Sec. 10. However, no matter what unfolding procedure is used, the same relation between reconstructed measurement $\vec{M}$ and truth distribution $\vec{S}$ given in Eqn. 9.1 is present, and must be accounted for through the estimation of the detector response R. In both forward folding and unfolding, this detector response is estimated through simulation, where the underlying truth values generated in the simulation are compared with the reconstructed event distribution. Therefore, no matter what type of analysis is being performed, it is important that the model prediction is accurate.

In many cases, the model prediction can be directly tested. For example, model predictions on muon kinematics can be directly compared to the observed distributions of muon kinematics seen in data. Furthermore, muons are both easily reconstructed as a result of their simple topology, and well understood from previous experiments and theory 5. These advantages give confidence in the validation of the modeling of reconstructed muon kinematics. By comparison, quantities such as neutrino energy cannot be tested so directly and thoroughly, and therefore their modeling must be considered very carefully in any analysis

that relies heavily on it. The neutrino energy cannot be observed directly like the muon energy, so it must be reconstructed from the final state particles following an interaction. However, in many detectors, some of these particles escape detection, such as neutrons in LArTPC detectors. To err on the side of caution, cross section analyses can avoid producing measurements over neutrino energy, and instead focus on fully observable quantities such as muon kinematics. However, oscillation analyses are afforded no such luxury as neutrino oscillations depend explicitly on neutrino energy, motivating the measurement of neutrino energy in cross section analyses (which serve to improve modeling in oscillation and other other physics search analyses) whenever its modeling can be trusted. Finally, even cross section analyses over observable variables rely on modeling to account for the biases of selection efficiency and purity, meaning that accurate modeling remains a key consideration of every analysis.

The goal of model validation is to ensure that the model and its uncertainties are capable of describing the distribution in data. As has been discussed in Sec. 8, model predictions cover an immense number of parameters, and even then they vastly over-simplify reality. As such, the data in its totality can be thought of as a massively higher-dimensional distribution. In practice, the data and MC prediction are always seen through reconstructed distributions over a few dimensions at most. This represents a marginalization of the higher dimensional distribution, where all non-present dimensions have been integrated over, hiding an data/MC differences. As mentioned earlier, the existence of of invisible detector activity further complicates attempts to validate model predictions in data.

Faced with these complications, many analyses use supplemental model predictions to test the coverage of the analysis model prediction. By using different model predictions, such as Neut [160], NuWro [161], and GiBUU [162], the complex phase space of possible model predictions can be explored in a physically motivated manner (to the extent that each model prediction is a reasonable description of the data). While the data distribution

over the full higher-dimensional space cannot be directly tested, nor can invisible activity in the detector, comparing model predictions can give a sense of the expected range of variations possible in data. A common stance is to require the analysis model prediction (including uncertainties) to cover the variation seen in other model predictions. Minerva considers the unfolded measurement when separately unfolded using their latest analysis model MnvTune-v3 (based on a modified version of Genie v2.8.4) and the earlier MnvTune-v1.2, and takes the difference between unfolded measurements as an additional uncertainty on the measurement [163, 164]. T2K unfolds a large number of fake data set generated from different models, and requires that the unfolded fake data “measurements" agree with the analysis prediction within half of the listed uncertainty. Both of these approaches ensure that any data mis-modelings of the analysis model that are within the scope of the other models considered would not bias the unfolded measurement beyond the listed uncertainties. However, these methods leave open the question of whether the set of physicist-produced models has sufficient variability to cover nature. Furthermore, in theory if a particular model does not describe data well, an analysis model shouldn't need to have enough uncertainty to cover it.

9.1 The Wire-Cell Model Validation Procedure

Wire-Cell adds the to the field of model validation techniques by introducing a procedure for data-driven validation [137] that allows for increased sensitivity to mis-modeling errors. This is made possible through the wealth of information available in the LArTPC detector and preserved through the Wire-Cell reconstruction, resulting in the non-zero efficiency throughout the kinematic phase space 64 that allows for a complete data/MC comparison. By directly evaluating model performance on data, this approach avoids having to determine whether comparisons to alternate models provide sufficient and reasonable coverage. It is possible that reality differs from the range of available model predictions, or that an alternate

model CV prediction does not well describe reality, in which case any tension between the analysis model and it should not give concern. Instead of relying on the accuracy of alternate models, the Wire-Cell validation procedure is able to leverage simultaneous measurements over different kinematic distributions to produce a stringent test of the capability to describe the distribution seen in data.

The key insight leveraged is that model predictions over different reconstructed variables are correlated, allowing measurements over well-reconstructed variables such as muon kinematics to narrow the model prediction over other distributions through the conditional constraint procedure. This creates a more stringent model test that is more sensitive to mis-modeling than the unfolded measurement, warning against unfolding in cases where there is significant model error. Furthermore, the correlated predictions across reconstructed variables can be simultaneously leveraged to construct a test that is sensitive to the modeling of invisible detector activity, which will be discussed in further detail in the following paragraphs. It is important to note that the conditional constraint procedure is only used to produce updated model predictions for validation purposes; unfolded measurements are performed with the default MicroBooNE model without applying any constraint.

The conditional constraint formalism uses a Bayesian approach to update the model prediction based on the data measurement over a separate distribution, in a format very similar to the GPR smoothing procedure used in Sec. 8.4.2. It starts by viewing the model prediction as a joint distribution over multiple variables of interest $\{a, b, ...\}$, which is easily achieved since the analysis arranges uncertainties using the covariance matrix formalism [141]:

$$p(a,b) = N\left(\begin{bmatrix}\mu_a \\ \mu_b\end{bmatrix}, \begin{bmatrix}\Sigma_{aa} & \Sigma_{ab} \\ \Sigma_{ba} & \Sigma_{bb}\end{bmatrix}\right) \tag{9.2}$$

The prior is taken to be the default model prediction, and a posterior distribution on a is

computed following Bayes theorem for a given measurement y_b:

$$\begin{aligned}\hat{p}(a|y_b) &= N\left(\hat{\mu}_{a|y_b}, \hat{\Sigma}_{aa|y_b}\right) \\ \hat{\mu}_{a|y_b} &= \mu_a + \Sigma_{ab}\left(\Sigma_{bb}\right)^{-1}\left(y_b - \mu_b\right) \\ \hat{\Sigma}_{aa|y_b} &= \Sigma_{aa} - \Sigma_{ab}\left(\Sigma_{bb}\right)^{-1}\Sigma_{ba}\end{aligned} \tag{9.3}$$

From the posterior prediction, the updated model prediction on a can be tested through a Pearson χ^2 Goodness of Fit (GoF) test between measurement $M = y_a$ and prediction $P = \hat{p}(a|y_b)$:

$$\chi^2 = (M-P)^T \times C^{-1} \times (M-P) \tag{9.4}$$

where C represents the full covariance matrix. Note that while the Pearson construction is more accurate for the GoF validation tests, the Combined Neyman-Pearson construction gives smaller bias and is used in the unfolded measurement [165, 166]. Under the null hypothesis, the model prediction is accurate, and the distribution in data should represent a reasonable draw from the posterior. However, if the model prediction is insufficient to describe the distribution in data, then a $\chi^2/ndf > 1$ and p-value< 0.05 will be observed, rejecting the null hypothesis. Note that in the case of $\chi^2/ndf << 1$, strictly speaking the data does not appear to be reasonable draw from the model; however, this is indicative of an over-estimation of uncertainties in the model, which suggests a conservative unfolded measurement but not significant bias from the model.

By applying the measurement on y_b as a constraint, the updated prediction on a is restricted, based on the correlated predictions over a and b, as shown in Fig. 83. Conceptually, this can be thought of as restricting the model parameters to ranges that are consistent with y_b. For example, an overall rate abundance observed in y_b may restrict the flux normalization parameter to only allow higher values. In this way, the uncertainties in common between a and b are significantly, reduced through the application of the constraint. While a straightfor-

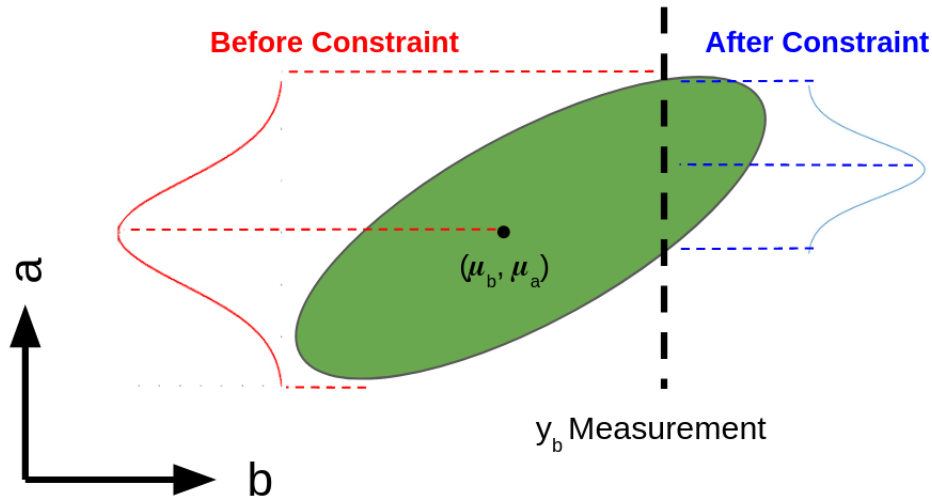

Figure 83: Illustration of conditional constraint procedure being used to update the model prediction over a given the measurement of y_b. The marginalised prior (posterior) prediction over a is shown in red (blue).

ward model comparison may fail to detect model deficiencies, owing to the combination of uncertainties across model components over the marginalized reco-space distribution, the conditional constraint procedure allows uncertainties of interest to be somewhat isolated. The resulting GoF test produces a focused evaluation of whether the remaining uncertainties can adequately cover the distribution seen in data over the target variable(s).

To help demonstrate the capabilities of the constrained GoF test, a series of fake data sets are used. In particular, the Genie v2.12.10 CV prediction is used to generate a set of fake data that is statistically independent from the MicroBooNE model prediction using a tuned version of Genie 3.0.6 (see Sec. 8.3). Since the Genie v2.12.10 CV prediction is known to disagree with the MicroBooNE model (and also with data), we expect to find tension in the GoF tests. Another fake data set is constructed from the MicroBooNE model prediction with artificially lowered proton energy, mimicking the possibility that the division of transfer energy ν between visible hadronic and missing energy is mis-modeled. In both cases, the goal of the GoF test is to identify the known mis-modeling, even when the bias introduced is not so severe as to exceed the listed uncertainties in the measurement. This is made possible

through the reduced uncertainties and focused model prediction in the constrained GOF test, allowing it to be more sensitive to mis-modeling than the unfolded data/MC χ^2 comparison. The performance of these fake data studies is discussed in more detail in Sec. 9.5.

9.2 Validation of the Modeling of Kinematic Event Distributions

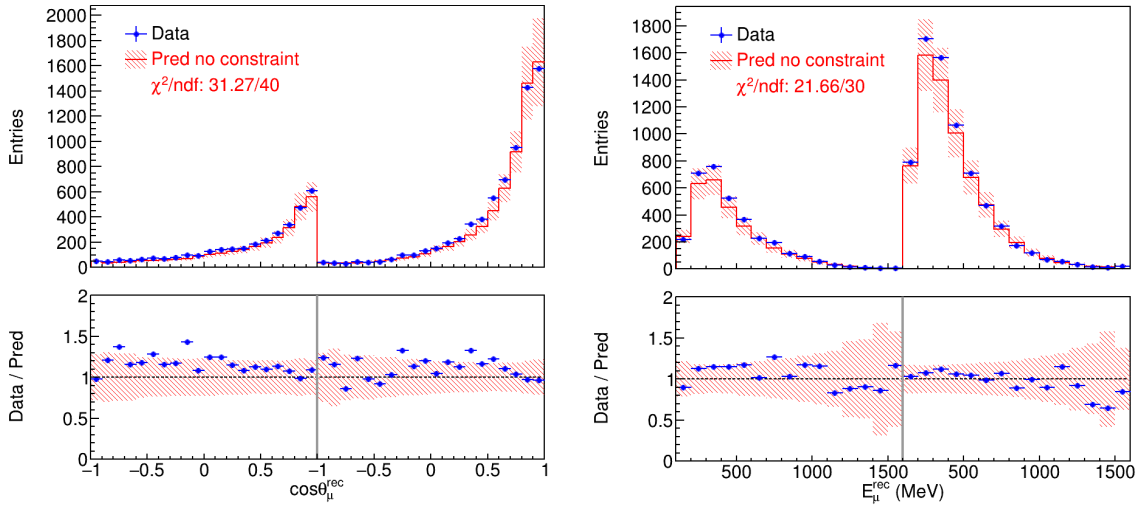

Figure 84: Left: Comparison between measured and predicted distributions of selected events over $cos(\theta_\mu)$. FC events are shown in the first 20 bins on the left and PC events are shown in the next 20 bins on the right. Right: Comparison between measured and predicted distributions of selected events over E_μ. FC events are shown from 0.1 GeV to 2.5 GeV in the first 15 bins (including overflow) on the left and PC events are shown in the next 15 bins on the right. The data/MC ratios are shown below the overall event distribution.

In this section, each kinematic variable used is evaluated under a standard GoF test, without applying the conditional constraint procedure. Later sections will investigate key distributions featuring difficult modelings in more detail by employing the conditional constraint. In each GoF test, a $\chi^2/ndf < 1$ and corresponding p-value< 0.05. indicate that the model is capable of describing the distribution seen in data. Fig. 84 shows the GoF tests over $cos(\theta_\mu)$ and E_μ, and Fig. 85 shows the GoF tests over E_ν and E_{had}^{reco}. Since 1D distributions probe a smaller section of phase space, the 2D distribution of $\{P_\mu, cos(\theta_\mu)\}$ is studied in Fig. 86 to test the multi-variable phase space used in the triple-differential measurement. 3D

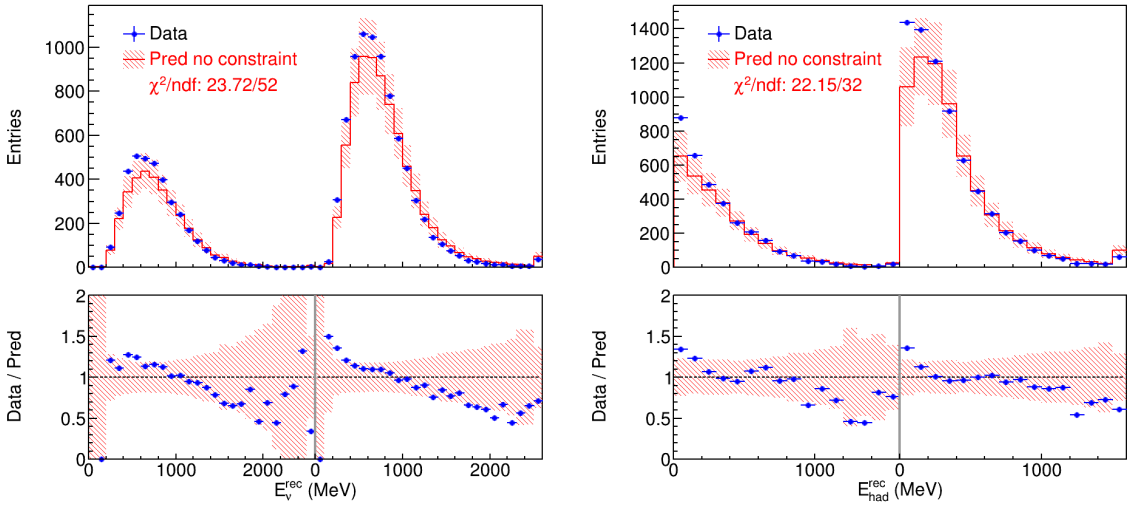

Figure 85: Left: Comparison between measured and predicted distributions of selected events over E_ν. FC events are shown from 0 GeV to 4 GeV in the first 26 bins (including overflow) on the left and PC events are shown in the next 26 bins on the right. Right: Comparison between measured and predicted distributions of selected events over E_{had}. FC events are shown from 0 GeV to 2.5 GeV in the first 16 bins (including overflow) on the left and PC events are shown in the next 16 bins on the right. The data/MC ratios are shown below the overall event distribution.

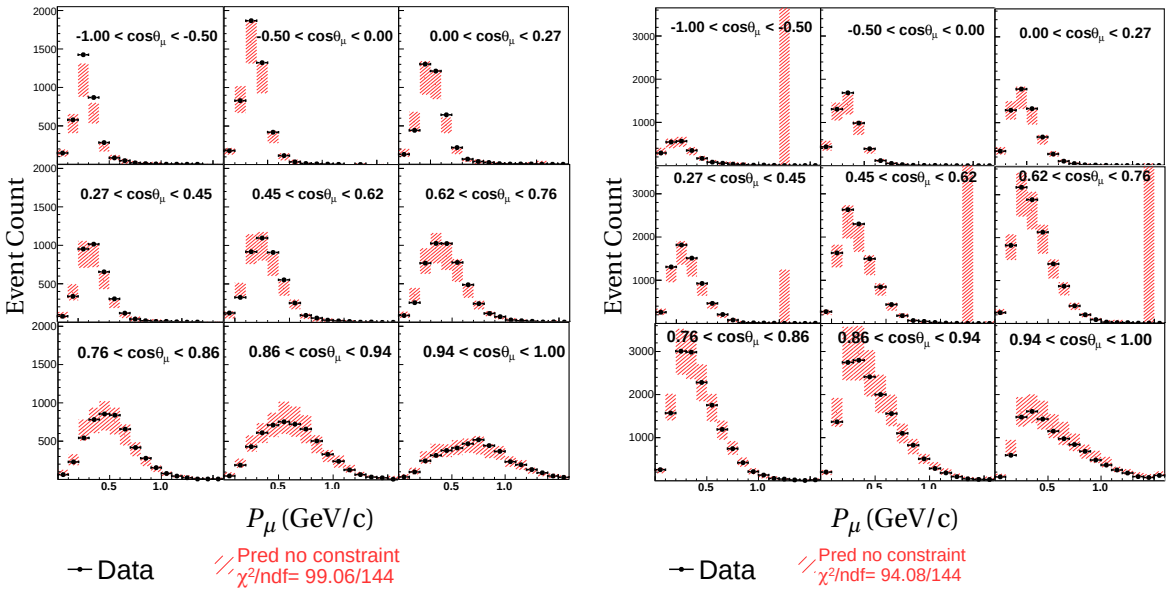

Figure 86: Comparison between measured and predicted distributions of selected FC (left) and PC (right) events over P_μ within each $cos(\theta_\mu)$ slice. The particularly large uncertainty in a few bins is because of very low statistics.

distributions are not used as they are over-constrained through the conditional constraint procedure given that the underlying kinematics are largely defined by three degrees of freedom. In the upcoming sections, 2D distributions under constraint will be studied, providing more sensitive tests of the model than a simple 3D distribution could.

9.3 Validation of the Modeling of Partially Contained Event Reconstruction

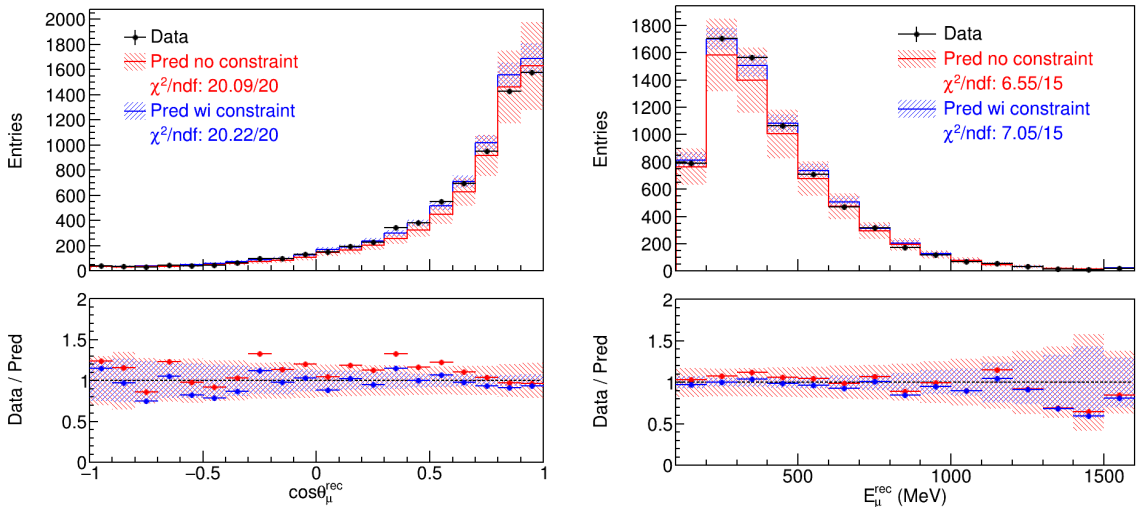

Figure 87: Left: Comparison between measured and predicted distributions of selected PC events over $cos(\theta_\mu)$ before (red) and after (blue) applying the measured FC distribution as a constraint. Right: Comparison between measured and predicted distributions of selected PC events over E_μ before (red) and after (blue) applying the measured FC distribution as a constraint. The data/MC ratios are shown below the overall event distribution.

PC events include activity outside the FV, making them a situation where invisible energy must be accounted for through modeling. This is particularly the case for muons, whose long tracks can extend a significant distance outside the detector volume. As mentioned earlier, 1D and 2D marginalizations of higher-dimensional distributions can hide mis-modeling through the collapse of various systematic uncertainties, many not related to the modeling of particle energies in PC events. Therefore, the (well understood) FC event distribution is used

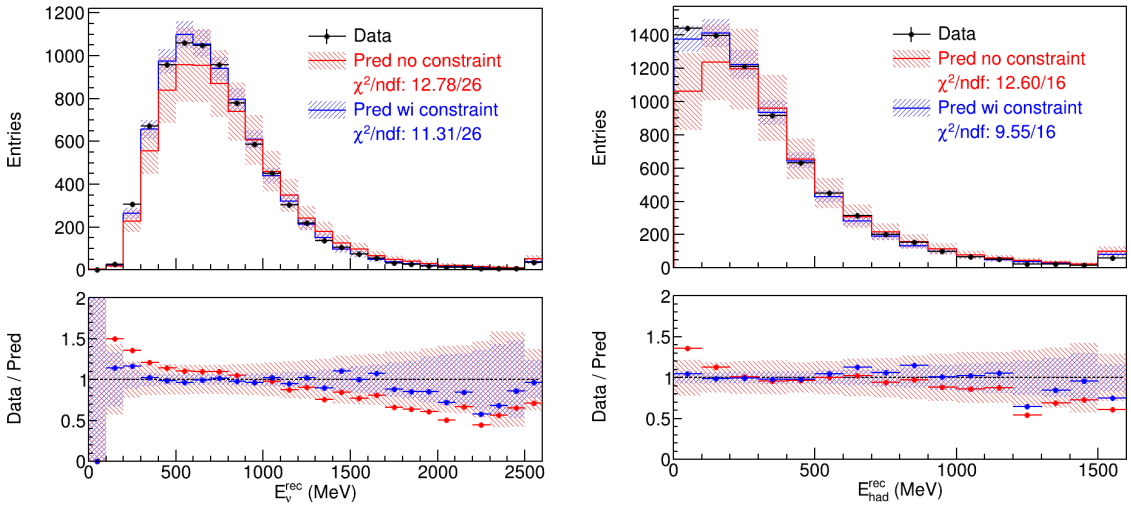

Figure 88: Left: Comparison between measured and predicted distributions of selected PC events over E_ν before (red) and after (blue) applying the measured FC distribution as a constraint. Right: Comparison between measured and predicted distributions of selected PC events over E_{had} before (red) and after (blue) applying the measured FC distribution as a constraint. The data/MC ratios are shown below the overall event distribution.

as a constraint on the model prediction of the PC event distribution to validate the mapping between truth and reconstructed PC event kinematics.

Fig. 87 shows the GoF tests over $cos(\theta_\mu)$ and E_μ for PC events, with the respective FC event distributions used as a constraint. Fig. 88 shows the GoF tests over E_ν and E_{had}^{reco} for PC events, with the respective FC event distributions used as a constraint. The red bands show the uncertainties before constraint, and the blue bands show the updated model prediction and uncertainties after constraint. Fig. 89 shows the Gof test over the 2D distribution of $\{P_\mu, cos(\theta_\mu)\}$ for PC events, with the FC event distribution used as a constraint. In each case, the model uncertainty is significantly reduced by the constraint, but the model prediction CV is also updated to a more accurate distribution, allowing the constrained model prediction to still cover the distribution in data. If there was significant mis-modeling in the PC events not seen in FC events, then applying the FC selection as a constraint would not improve the model prediction CV, causing the constrained GoF test χ^2/ndf to increase significantly. This was not seen in the data, and instead each GoF test shows good data/MC agreement even

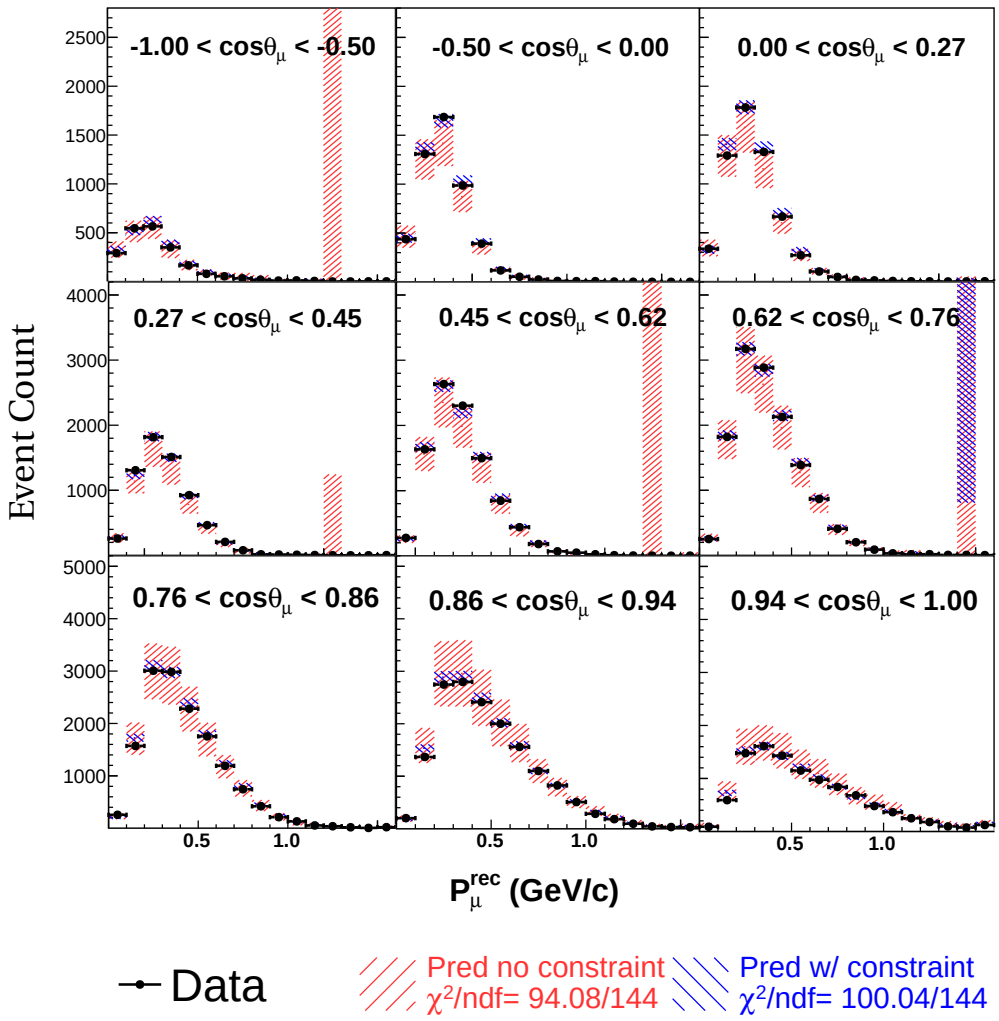

Figure 89: Comparison between measured and predicted distributions of selected PC events over P_μ within each $cos(\theta_\mu)$ slice before (red) and after (blue) applying the measurement of the distribution of FC events over $\{P_\mu, cos(\theta_\mu)\}$. The particularly large uncertainty in a few bins is because of very low statistics.

after applying the constraint.

9.4 Validation of the Modeling of Transfer Energy

The modeling of the energy transferred to the argon system, ν, is critically important to neutrino experiments. As has been discussed already in Sec. 2.1 and Sec. 9, accurate neutrino energy reconstruction is important for both oscillation analyses and cross section analyses,

and the transfer energy represents one of two components in the neutrino energy: $E_\nu = E_\mu + \nu$. Furthermore, ν includes energy carried away in non-ionizing particles such as neutrons and low-energy photons, making it a primary source of invisible activity in LArTPC detectors. This is described by separating the transfer energy into visible and missing components: $\nu = E_{had}^{vis} + E_{had}^{missing}$, where the reconstructed quantity E_{had}^{reco} serves as an estimate of E_{had}^{vis}, up to inaccuracies in the reconstruction and event selection. Therefore, the mapping from E_{had}^{reco} to ν and ultimately to E_ν relies on accurate modeling, and in particular cross section modeling. Reliable model validation is necessary for unfolding to E_ν, or else mis-modelings in the the prediction of ν may introduce significant bias into the unfolded measurement.

While $E_{had}^{missing}$ cannot be directly measured, the GoF test on the distribution of E_{had}^{reco} using the measurement of E_μ as a constraint is sensitive to the modeling of $E_{had}^{missing}$, because of the correlated predictions of muon kinematics and hadronic energy. There are two intuitive arguments to help see this fact. First is the argument of conservation of energy: $E_\nu = E_\mu + \nu = E_\mu + E_{had}^{vis} + E_{had}^{missing}$. The constrained GoF test directly leverages measurements of E_μ (through P_μ) and E_{had}^{vis} (through E_{had}^{reco}). Furthermore, the accurate measurement of the distribution over E_μ creates a constraint on the flux modeling parameters, and therefore on the E_ν prediction. This leaves $E_{had}^{missing} = E_\nu - E_\mu - E_{had}^{vis}$ as the only unknown to be solved for, making the constrained GoF test sensitive to it through the simultaneous measurements of the distributions over P_μ and E_{had}^{reco}.

Another way to view the sensitivity to the modeling of missing energy is to consider the underlying interaction channels that comprise the model prediction, and how they are updated through the application of the constraint. Model predictions over P_μ, E_{had}^{reco}, ... are comprised of predictions on the QE, RES, MEC, and DIS interactions, each withe their own distribution over each kinematic variable. By using the measured muon kinematic distribution as a constraint, the model predictions on the interaction channels are updated akin to a constrained re-weighting. Then, the updated QE, RES, MEC, and DIS predictions

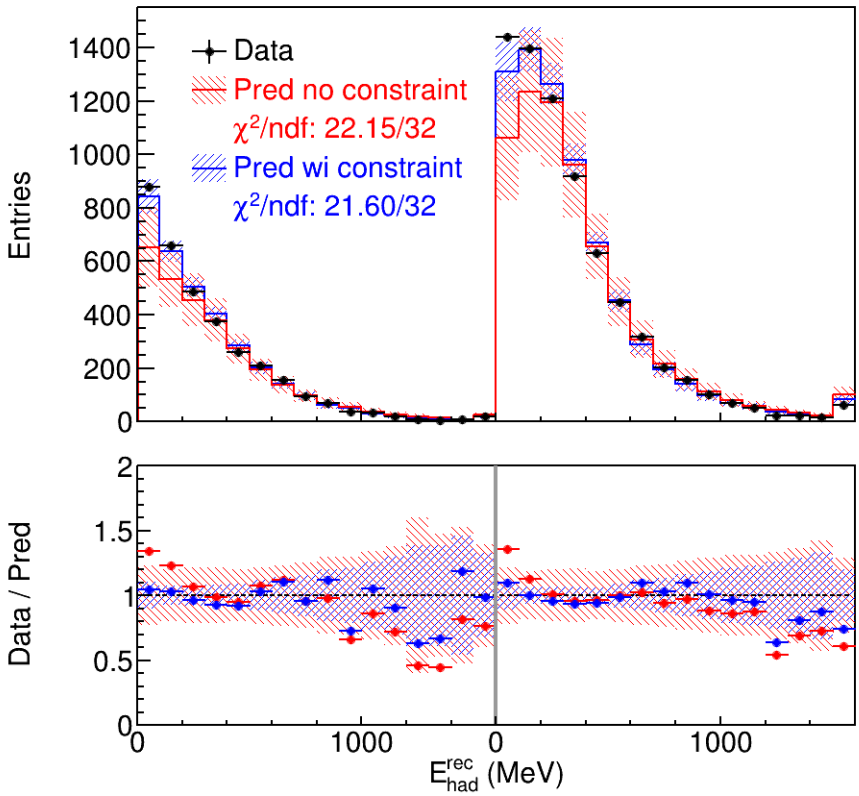

Figure 90: Comparison between measured and predicted distributions of selected events over E_{had}^{reco} before (red) and after (blue) applying the measured E_{μ} distribution as a constraint. FC events are shown from 0 GeV to 2.5 GeV in the first 16 bins (including overflow) on the left and PC events are shown in the next 16 bins on the right. The data/MC ratios are shown below the overall event distribution.

create new predictions over E_{had}^{reco} and $E_{had}^{missing}$. If there is mis-modeling in $E_{had}^{missing}$, then there should also be mis-modeling in E_{had}^{reco} that the GoF test is sensitive to because of the underlying physics in common. As a result, the constrained GoF test is capable of validating the modeling of correlations between muon kinematics and hadronic energy, and is sensitive to the modeling of missing energy.

Fig. 90 shows the GoF test over E_{had}^{reco}, using the measurement of the distribution over E_{μ} as a constraint. Fig. 91 shows the GoF test over $\{E_{had}^{reco}, cos(\theta_{\mu})\}$, using the measurement of the distribution over $\{P_{\mu}, cos(\theta_{\mu})\}$ as a constraint. In each case, the constrained χ^2 remains roughly constant despite the significant reduction in uncertainties, demonstrating the accurate modeling of correlations between the constraint and target distributions. The

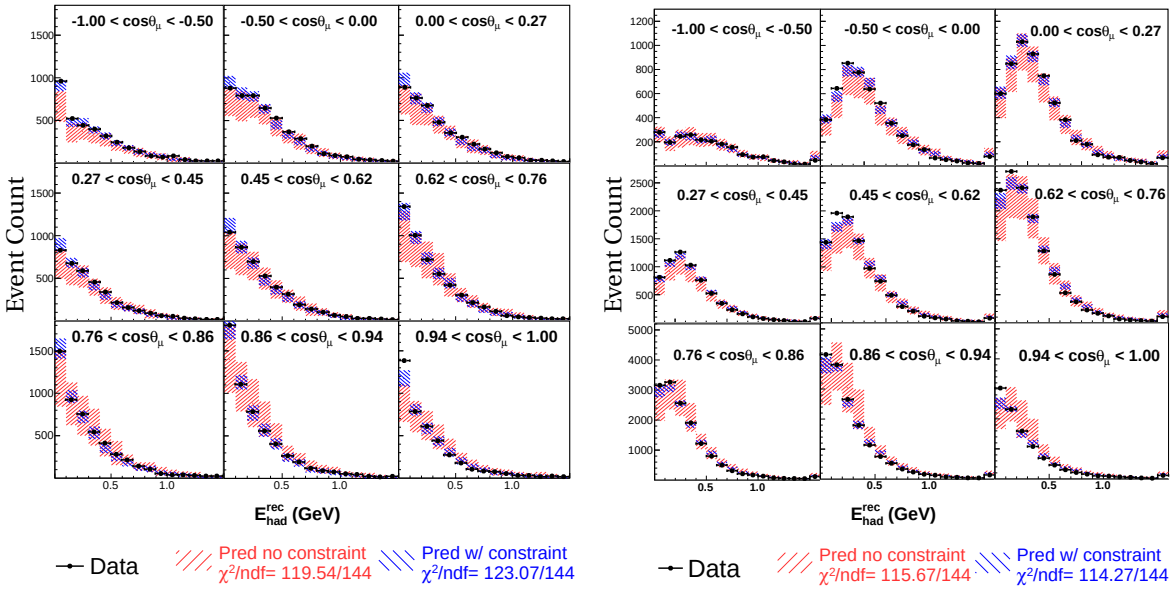

Figure 91: Comparison between measured and predicted distributions of selected FC (left) and PC (right) events over E_{had}^{reco} within each $cos(\theta_\mu)$ slice. The model prediction before (after) applying the measurement of the distribution over $\{P_\mu, cos(\theta_\mu)\}$ is shown in red (blue).

successful constrained GoF tests over E_{had}^{reco} give confidence in the mapping from E_{had}^{reco} to ν, and therefore the modeling of E_ν within uncertainties. Any bias from mis-modeling of ν introduced through unfolding should be within the listed uncertainties of the unfolded measurement.

9.5 Examining the Goodness of Fit Test with Fake Data

The modeling of ν discussed in Sec. 9.4 takes the focus of these fake data studies, as it is the largest source of invisible activity, and most prone to mis-modeling. Fake data allows the reasoning presented in the previous sections to be tested by creating conditions where we may expect model failure and observing how the GoF tests perform. Specifically, the constrained GoF tests should be more stringent than the overall unfolded measurement, meaning that it should be possible to demonstrate a GoF test that detects mis-modeling even when the bias introduced in the unfolded measurement is within uncertainty. This level of forewarning gives significant confidence to the unfolding and overall measurement, seeing as

the GoF tests serve as the canary in the coal mine - detecting problems before the disastrously harm the overall measurement.

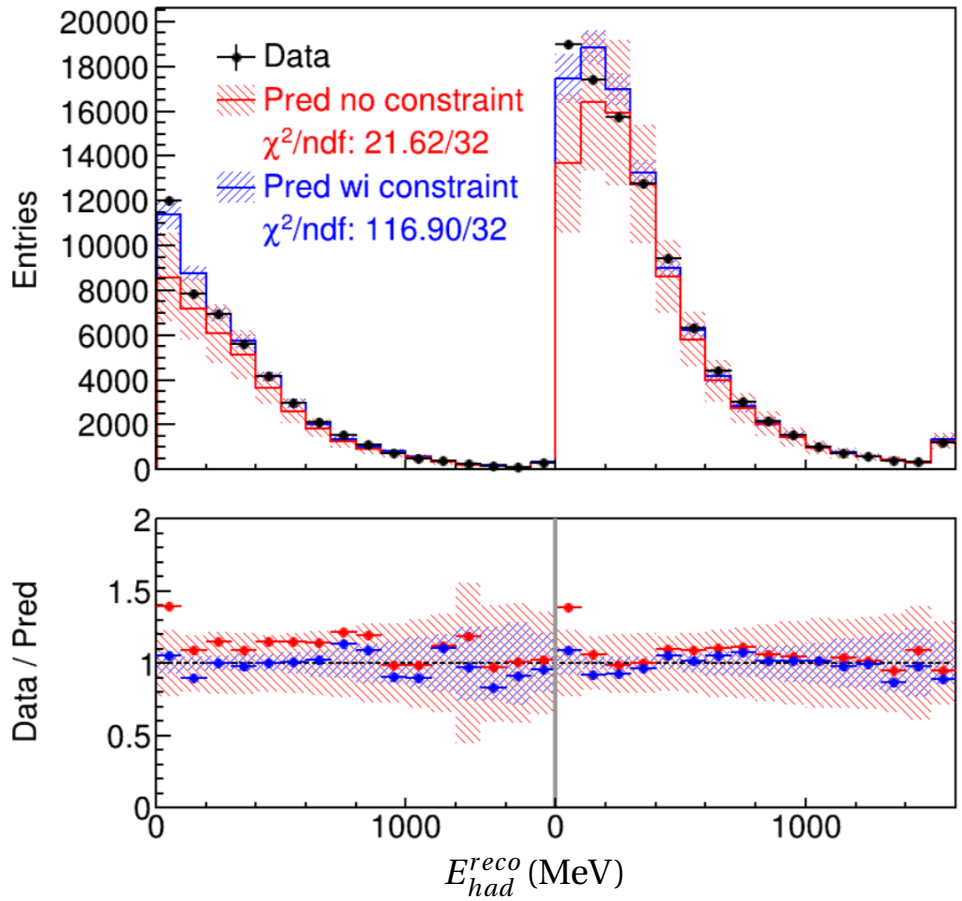

Figure 92: Comparison between measured and predicted distributions of selected events over E^{reco}_{had} before (red) and after (blue) applying the measured $\{E_\mu, cos(\theta_\mu)\}$ distribution as a constraint. FC events are shown from 0 GeV to 2.5 GeV in the first 16 bins (including overflow) on the left and PC events are shown in the next 16 bins on the right. The data/MC ratios are shown below the overall event distribution. Source: [167]

The first fake data set, consisting of 7.2×10^{20} POT, is generated using Genie v2.12.10, which is known to disagree with the MicroBooNE model based on Genie v3.0.6 [167]. Note that as an entirely separate set of MC events, this fake data set is fully statistically independent from the MicroBooNE model. Fig. 92 shows the comparison between measured and predicted distributions of events over E^{reco}_{had} both before and after applying the measurement of the distribution of events over $\{E_\mu, cos(\theta_\mu)\}$ as a constraint. The unconstrained model prediction

is unable to detect mis-modeling, owing to the fact that it is a 1D marginalization of higher-dimensional distribution. However, after applying the constraint, the tension between fake data and MicroBooNE model prediction grows significantly, rising to $\chi^2/ndf = 16.9/32$ with a corresponding p-value of 1.338×10^{-11}. This demonstrates the increased sensitivity of the constrained GoF test, which only becomes apparent in situations with significant tension between measurement and prediction.

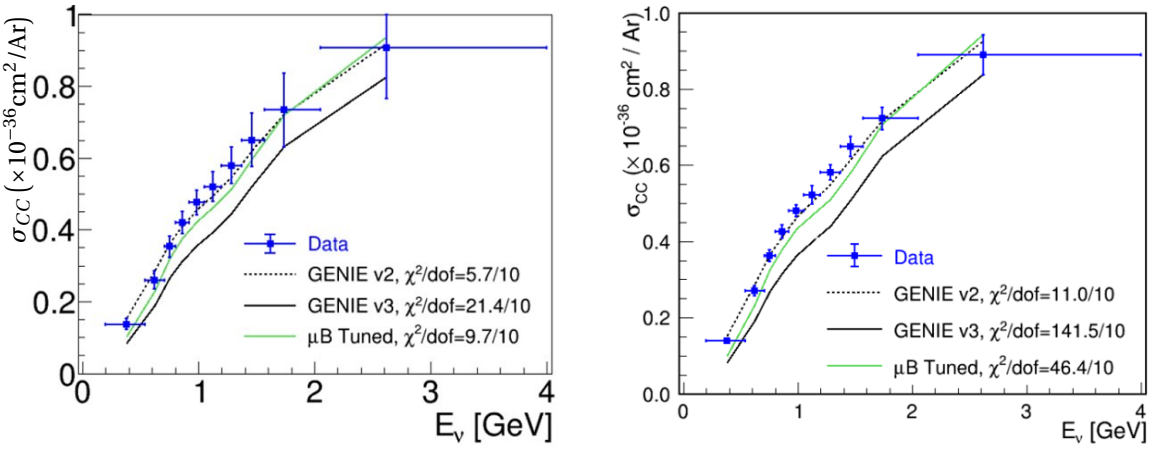

Figure 93: Extracted total cross section as a function of E_ν, using Genie v2.12.10 as fake data. Left: Unfolded with the full covariance matrix. Right: Unfolded with only cross section, stat, and MC stat uncertainties for a more accurate comparison, given that the fake data only deviates in the cross section modeling. Three model predictions, including GENIE v2, GENIE v3 default tune, and Genie v3 MicroBooNE-tune, are shown. Source: [167]

The results of the GoF test would warn against unfolding in the case of real data; however, we can proceed with unfolding to judge the level of bias introduced by the mis-modeling that was detected. Fig. 93 shows the unfolded total cross section measurement as a function of neutrino energy, considered separately when unfolded using the full covariance matrix and a modified version with only cross section and statistical uncertainties. Since the model differences between the fake data and the MicroBooNE model are limited to the cross section prediction (and not, say, flux), the limited covariance matrix gives a more accurate comparison. In both cases the bias introduced is not nearly as large as the disagreement detected in the GoF test, with χ^2/ndfs of 5.7/10 and 11/10 for the full and limited covariance matrices,

Table 4: χ^2 and p-value of constrained E_{had}^{reco} GoF test using proton-energy-scaled fake data. Only stats, MC stats, and cross section uncertainties were used. Source: [167]

E_p^{reco} Scaling Factor	$\chi^2(ndf = 32)$	p-value
0.95	5.34	1.0
0.9	21.05	0.93
0.85	47.01	0.04
0.8	80.6	0.00

and p-values of 0.84 and 0.36 respectively. This fake data study exemplifies the capability of the constrained GoF test to detect mis-modeling with more sensitivity than the unfolded measurement.

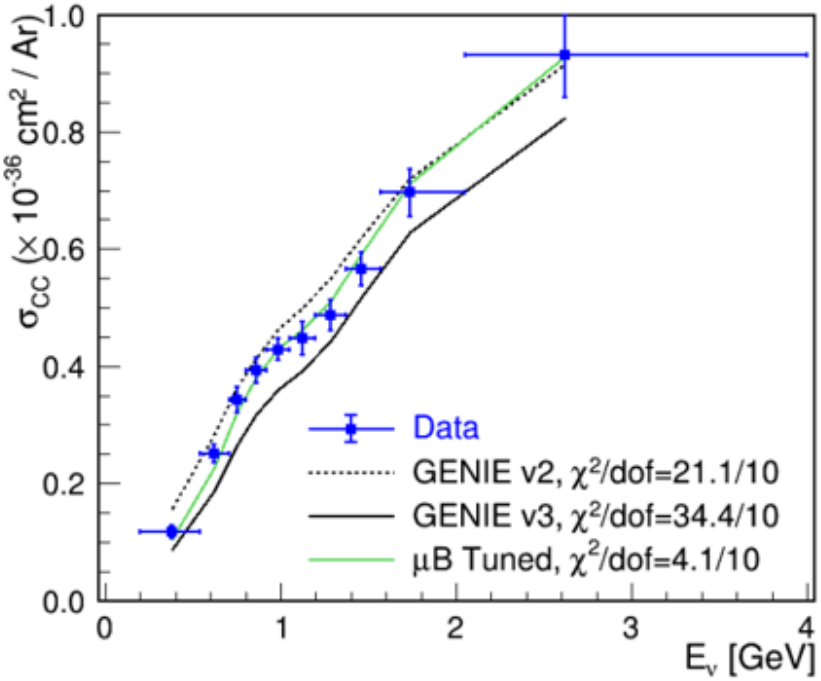

Figure 94: Extracted total cross section as a function of E_ν, using the MicroBooNE MC with 85% proton energy scaling as fake data. The measurement was unfolded with only cross section, stat, and MC stat uncertainties for a more accurate comparison, given that the fake data only deviates in the cross section modeling. Three model predictions, including GENIE v2, GENIE v3 default tune, and Genie v3 MicroBooNE-tune, are shown. Source: [167]

An additional fake data set is constructed to specifically test the hypothesis of mis-

modeling the division of transfer energy between visible and missing components. This data set is generated by taking the MicroBooNE model MC and scaling down the reconstructed proton energy in each interaction, simulating a case where the lost energy was transferred away in neutrons. Proton energies were scaled down by 5%, 10%, 15%, and 20%, and GoF tests performed in each case, using only statistical and cross section uncertainties. The results of each GoF test is shown in Table 4. At 15% proton energy loss the GoF test is able to detect the mis-modeling with $\chi^2/ndf = 47.01/32$ and a p-value of 0.04. In this scenario the cross section measurement was extracted, shown in Fig. 94. A $\chi^2/ndf = 4.1/10$ with the MicroBooNE model prediction indicates a low amount of bias introduced, demonstrating that the model validation procedure is able to detect mis-modeling in the transfer energy before it significantly biases the extracted measurement.

Since this data set is generated from the MicroBooNE MC events, they are not statistically independent. In theory, the extracted cross section should be unfolded without statistical uncertainties. However, this presents a practical in-feasibility, as the statistical uncertainties give the covariance matrix diagonal terms that allow a pseudo-inverse to be computed. Still, the MC data set is large, at roughly 10 times the size of the real data, so statistical fluctuations are small. Given these circumstances, the cross section is extracted using statistical and cross section uncertainties only, which is a reasonable approximation of cross section statistics only. Furthermore, the GoF test is constructed from statistical and cross section uncertainties, so if all statistical uncertainties were removed for a fully faithful treatment, both the extracted cross section bias and the model validation tension would grow, preserving the status-quo where the model validation is a more stringent test of mis-modeling.

10 Unfolding Procedure

As discussed in Sec. 9, unfolding is the process of solving for an underlying signal distribution s given a measurement m, background prediction b, and detector response r that smears the signal distribution, presented in Eqn. 9.1. For convenience, here the background-subtracted measurement $m' = m - b$ will be used to construct the master equation for unfolding:

$$m' = r \cdot s \tag{10.1}$$

10.1 Unfolding without Regularization

The naive approach to unfolding is to simply solve for S by computing the pseudo-inverse of R and applying it to each side. Accounting for the uncertainty on M' in the form of the covariance matrix Cov, this is performed by minimizing the test statistic χ^2_0 describing the difference between background-subtracted measurement and smeared response:

$$\chi^2_0 = (m' - r \cdot s)^T Cov^{-1} (m' - r \cdot s) \tag{10.2}$$

Through Cholesky decomposition [168], the covariance matrix can be decomposed into triangular matrices and used to pre-scale the measurement and response matrix:

$$\begin{aligned} Cov^{-1} &= Q^T \cdot Q \\ M' &= Q \cdot m' \\ R &= Q \cdot r \end{aligned} \tag{10.3}$$

This leads to the simplified form of the test statistic:

$$\chi^2_0 = (M' - R \cdot s)^T (M' - R \cdot s)$$

$$= \sum_i \left(M'_i - \sum_j R_{ij} s_j \right)^2 \tag{10.4}$$

with solution $\hat{s}$ given by:

$$\hat{s} = \left(R^T R\right)^{-1} R^T \cdot M' \tag{10.5}$$

This solution can be simplified by considering the Singular Value Decomposition (SVD) of R:

$$R = U \cdot D \cdot V^T \tag{10.6}$$

which expresses R in terms of orthogonal matrices U and V that satisfy $UU^T = U^T U = I$ and $VV^T = V^T V = I$, as well as the diagonal matrix D composed of the singular values d_i arranged in descending order. The simplified solution for $\hat{s}$ becomes:

$$\begin{aligned} \hat{s} &= V \cdot D^{-1} \cdot U^T \cdot M' \\ &= V \cdot D^{-1} \cdot \cdot M'_U \end{aligned} \tag{10.7}$$

where $M'_U = U^T M'$ represents the background-subtracted measurement in the singular-value basis.

This naive solution is prone to instability resulting from measurement noise. This can be shown by expressing M' as a combination of the true underlying signal s_{true} and the measurement noise N:

$$\begin{aligned} M' &= R \cdot s_{true} + N \\ M'_U &= R_U \cdot s_{true} + N_U \\ \hat{s} &= V \cdot D^{-1} \cdot (R_U \cdot s_{true} + N_U) \end{aligned} \tag{10.8}$$

where $R_U = U^T R$ and $N_U = U^T N$ are the transformed response matrix and noise in the

singular-value basis, respectively. Notably, the noise terms are uncorrelated and all Gaussian distributed following $N(\vec{0}, I)$ (in the pre-scaled coordinate system), so after re-ordering by the orthogonal matrix U^T the noise terms are left unchanged. The least impactful eigenvectors in the SVD basis have correspondingly small singular values d_i, which causes the corresponding elements $1/d_i$ in D^{-1} to be very large. Normally, the corresponding smeared signal elements will be very small to counteract the influence of large elements of D^{-1}, however, the presence of noise allows for large fluctuations in the extracted signal.

10.2 Wiener SVD Unfolding

Typically this instability can be addressed through the use of regularization, which suppresses fluctuations through the inclusion of a penalty term λ in the χ^2 that constrains the extracted signal distribution:

$$\chi^2 = \chi_0^2 + \lambda \tag{10.9}$$

Often the smoothness of the signal distribution is asserted by constructing the penalty term from a derivative (1st, 2nd, 3rd, ...) of $\hat{s}$. The inclusion of a regularization term in the test statistic smears the extracted signal distribution, described through the additional smearing matrix A:

$$A = V \cdot F \cdot V^T \tag{10.10}$$

$$\begin{aligned} \hat{s} &= A \cdot V \cdot D^{-1} \cdot M'_U \\ &= V \cdot F \cdot D^{-1} \cdot M'_U \end{aligned} \tag{10.11}$$

for some matrix F in the SVD basis.

The goal of regularization is to suppress the eigenvectors in SVD that contain high noise relative to the signal while leaving signal-dominant modes undisturbed. The Wiener Filter [169, 104], used earlier in Sec. 5.1.2, is designed to do just this, by suppressing eigenvectors

based on their signal-to-noise ratio. By substituting the expectation value of the unknown true signal distribution $\overline{s}_{true}$ with the expectation value of the extracted signal distribution $\overline{s}$ and using Eqn. 10.6, the terms of the Wiener filter can be computed for each SVD bin i:

$$\overline{R_U^2 \cdot s^2} = D \cdot V^T \cdot \overline{s^2} = d_i^2 \left(\sum_j V_{ij}^T \cdot \overline{s}_j \right)^2$$
$$\overline{N^2} = 1 \tag{10.12}$$

The Wiener filter is constructed as:

$$W_{ik} = \frac{\overline{R_U^2 \cdot s^2}}{\overline{R_U^2 \cdot s^2} + \overline{N^2}} = \frac{d_i^2 \cdot \left(\sum_j V_{ij}^T \cdot \overline{s}_j \right)^2}{d_i^2 \cdot \left(\sum_j V_{ij}^T \cdot \overline{s}_j \right)^2 + 1} \tag{10.13}$$

This matrix takes the place of F in determining the regularization and subsequent additional smearing A.

To better apply the intuitive goal of smoothness, the Wiener filter can be modified to regularize the curvature of the spectrum rather than simply the strength. This is achieved by introducing the matrix C, commonly taken to be C_1 or C_2, which take the first and second derivative respectively, into the definitions of R_U and $\overline{s}$:

$$\overline{R_U \cdot s} = (R_U \cdot C^{-1}) \cdot (C \cdot \overline{s}) \tag{10.14}$$

In this triple-differential analysis, C_3 representing the third derivative is chosen as it introduces the least bias into the extracted distribution. This is because most slices of the data form the shape of a downward parabola, representing a significant second derivative. Applying a penalty term based on the first or second derivative would then attempt to flatten the distribution in a consistently biased manner. However, the general first-order shape of the data presents no significant third derivative, reducing the shape bias introduced As will

be discussed in more detail shortly, it is worth noting that no matter what regularization is chosen within the Wiener SVD method, so long as the additional smearing matrix A_C is computed the bias introduced is known, and can be applied to any model comparison to account for its effect, akin to forward folding.

By redefining the response matrix, the SVD decomposition basis is changed, adjusting the computation of the extracted signal, additional smearing matrix, and Wiener filter:

$$R \cdot C^{-1} = U_C \cdot D_C \cdot V_C^T \tag{10.15}$$

$$\hat{s} = A_C \cdot (R^T R)^{-1} \cdot R \cdot M_U \tag{10.16}$$

$$A_C = C^{-1} \cdot V_C \cdot W \cdot V_C^T \cdot C \tag{10.17}$$

$$W_{ik} = \frac{d_{Ci}^2 \cdot \left(\sum_j V_{Cij}^T \cdot \left(\sum_l C_{jl} \cdot \bar{s}_l \right) \right)^2}{d_{Ci}^2 \cdot \left(\sum_j V_{Cij}^T \cdot \left(\sum_l C_{jl} \cdot \bar{s}_l \right) \right)^2 + 1} \tag{10.18}$$

In this manner, the Wiener filter can be customized to be applied to whatever derivative of the extracted signal distribution is wanted. Regularizing the derivative of the signal distribution proves to be an overall better way to control instability in unfolding [169]. The Wiener filter has the advantage of maximizing the overall signal-to-noise ratio in the effective frequency domain (SVD basis). This allows it to effectively control fit instability fluctuations without over-suppressing the desired signal, all without requiring the regularization strength to be tuned. Furthermore, the additional smearing introduced through the Wiener filter regularization is captured in the matrix A_C, which is computed and published alongside the extracted cross section measurement. In this way, any model can be directly compared to the measurement by simply applying the additional smearing matrix.

Since unfolding is a linear transformation from the measurement m over reconstructed variables to the signal distribution $\hat{s}$ over truth variables, it is possible to easily compute the covariance matrix in the unfolded domain. This is done by recognizing that Eqn. 10.16 is

essentially a relation between these two quantities:

$$\hat{s} = R_{tot} \cdot m \tag{10.19}$$

$$R_{tot} = A_C \cdot \left(R^T R\right)^{-1} \cdot R^T \cdot Q \tag{10.20}$$

This allows the unfolded covariance matrix to be easily deduced as:

$$Cov_{\hat{s}} = R_{tot} \cdot Cov \cdot R_{tot}^T \tag{10.21}$$

10.3 Regularization in Multiple Dimensions

The procedure outlined above relies on taking derivatives of the signal distribution $\hat{s}$. In a single dimension this is a simple procedure, and can simply be computed from the differences of adjacent bins. However, across a multi-dimensional distribution, taking a derivative requires defining a path. So as to not preference any particular dimension in the choice of path, a more complicated combined derivative is computed from a combination of simple derivatives. For each dimension, the derivative along a single variable is computed separately, computed using the third derivative matrices C_{3,E_ν}, C_{3,θ_μ}, and C_{3,P_μ}. Then a combined derivative matrix is computed by adding each C_3 matrix in quadrature:

$$\begin{aligned} C_{3,3D}^2 &= C_{3,E_\nu} \cdot C_{3,E_\nu}^T + C_{3,\theta_\mu} \cdot C_{3,\theta_\mu}^T + C_{3,P_\mu} \cdot C_{3,P_\mu}^T \\ C_{3,3D} &= \sqrt{C_{3,3D}^2} \end{aligned} \tag{10.22}$$

The square root of $C_{3,3D}^2$ can be computed by diagonalizing it into VDV^T, which is guaranteed to be possible since since $C_{3,3D}^2$ is symmetric and positive definite, Then the square root of each diagonal element of D is taken, before recombining. The resultant matrix can be thought of as computing a mixed derivative that treats all three physical dimensions equally. This approach has been found to introduce less bias and reduce the instability fluctuations of the

unfolding when compared to simpler solutions such as simply using D_{P_μ} or $D_{cos(\theta_\mu)}$ as the derivative matrix of choice. Additionally, it provides a general procedure that can be applied to Wiener SVD unfolding in any multi-dimensional analysis, regardless of the number or content of variables.

It is worth noting that the distortion introduced in unfolding resulting from regularization as well as the Wiener filter are captured in the additional smearing matrix. This is applied to model predictions, so overall there is no change to the goodness of fit χ^2 for any model comparison. Still, smearing distortions can impact how we interpret the data, and should be reduced where possible. The Wiener filter inherently reduces the overall cross section measurement based on the signal to noise ratio, creating a smeared result with an overall normalization bias. A triple-differential cross section measurement contains a large number of bins, and so the Wiener filter can have a significant suppression effect on the unfolded cross section. To address this issue, the total 1-bin cross section prediction is compared before and after applying the additional smearing matrix, and a 21% deficit is observed. This 1.21 ratio is then applied to the data and MC as an overall rate re-normalization, effectively generating a re-normalized additional smearing matrix, leaving all model comparison χ^2 computations unchanged.

10.4 Estimation of the Detector Response Matrix

So far the discussion has assumed that the detector response R mapping the underlying signal s to the measurement m is known. To compute this, it is helpful to rewrite the measurement in terms of the various detector effects involved in producing it [167]:

$$m(E_\nu^{reco}) = POT \cdot T \int F(E_\nu) \cdot \sigma(E_\nu) \cdot D(E_\nu, E_\nu^{reco}) \cdot \epsilon(E_\nu, E_\nu^{reco}) \cdot dE_\nu + b(E_\nu^{reco}) \qquad (10.23)$$

where POT represents the number of protons on target, T represents the number of target nuclei, and E_ν and E_ν^{reco} represent the true and reconstructed neutrino energy, respectively.

$F(E_\nu)$ represents the neutrino flux as a function of neutrino energy, and $D(E_\nu, E_\nu^{reco})$ represents the detector response matrix, which describes the smearing from true to reconstructed neutrino energy. $\epsilon(E_\nu, E_\nu^{reco})$ represents the selection efficiency, and $b(E_\nu^{reco})$ represents the estimated number of background events. By grouping the various terms, Eqn. 10.23 can be simplified by expressing it in terms of the background events b_i and selected events $\tilde{s}_{ij}$ in each reconstructed bin i and truth bin j:

$$m_i = \sum_j \tilde{s}_{ij} + b_i \tag{10.24}$$

By grouping the terms in Eqn. 10.23 as well as employing some cancellation, $\tilde{s}_{ij}$ can be expressed in terms of the nominal-flux averaged total cross section s_j in each true neutrino energy bin, a flux constant $\tilde{F}_j$ that can be computed from the predicted flux, and the ratio of events selected in bin i that originate in bin j to the total number of events generated in bin j, $\tilde{\Delta}_{ij}$:

$$\begin{aligned}
s_j &= \frac{\int_j \overline{F}(E_{\nu,j}) \cdot \sigma(E_{\nu,j}) \cdot dE_{\nu,j}}{\int_j \overline{F}(E_{\nu,j}) \cdot dE_{\nu,j}} \\
\tilde{F}_j &= POT \cdot T \int_j \overline{F}(E_{\nu j}) \cdot dE_{\nu j} \\
\tilde{\Delta}_{ij} &= \frac{POT \cdot T \int_j F(E_{\nu,j}) \cdot \sigma(E_\nu, j) \cdot D(E_{\nu,j}, E_{\nu,i}^{reco}) \cdot \epsilon(E_{\nu,j}, E_{\nu,i}^{reco}) \cdot dE_{\nu,j}}{POT \cdot T \int_j \overline{F}(E_{\nu,j}) \cdot \sigma(E_\nu, j) \cdot dE_{\nu,j}} \\
\tilde{s}_{ij} &= \tilde{\Delta}_{ij} \cdot \tilde{F}_j \cdot s_j
\end{aligned} \tag{10.25}$$

where $\overline{F}$ represents the nominal (central value) muon neutrino flux. By using the nominal flux $\overline{F}$, we avoid having to consider the flux prediction uncertainties in the construction of the response matrix.

The above description can be generalized to a procedure to extract multi-differential cross sections, and in particular the triple-differential cross section in this analysis. In this case,

following the format of Eqn. 10.23, the measurement can be expressed as:

$$M\left(E_\nu^{reco},\theta_\mu^{reco},P_\mu^{reco}\right) = POT\cdot T\cdot\int\int\int F(E_\nu)\cdot\frac{d\sigma\left(E_\nu,\theta_\mu,P_\mu\right)}{d\theta_\mu dP_\mu}\cdot D\cdot\epsilon\cdot dE_\nu\cdot d\theta_\mu\cdot dP_\mu + B\left(E_{rec},\theta_\mu^{reco},P_\mu^{reco}\right) \tag{10.26}$$

where θ_μ and P_μ represent truth variables corresponding to the reconstructed muon scattering angle θ_μ^{reco} and muon momentum P_μ^{reco}, respectively. $\frac{d\sigma(E_\nu,\theta_\mu,P_\mu)}{d\theta_\mu dP_\mu}$ represents the differential cross section as a function of the truth variables. The detector response D and selection efficiency ϵ are both functions of all six truth and reconstructed variables. From this, a corollary to Eqn. 10.24 is formed for the 3D measurement bins by summing over reconstructed bins $\{i,j,k\}$ and truth bins $\{l,m,n\}$:

$$m_{ijk} = \sum_{lmn}\tilde{s}_{ijklmn} + b_{ijk} \tag{10.27}$$

In turn, the nominal-flux averaged differential cross section s_{lmn}, flux constant $\tilde{F}_{lmn}$, selection mapping $\tilde{\Delta}_{ijklmn}$, and selection count $\tilde{s}_{ijklmn}$ can be computed:

$$\begin{aligned}
s_{lmn} &= \frac{\int_{lmn}\overline{F}\left(E_{\nu,l}\right)\cdot\frac{d\sigma(E_{\nu,l},\theta_{\mu,m},P_{\mu,n})}{d\theta_{\mu,m}dP_{\mu,n}}\cdot dE_{\nu,l}\cdot d\theta_{\mu,m}\cdot dP_{\mu,n}}{\int_{lmn}\overline{F}\left(E_{\nu,l}\right)\cdot dE_{\nu,l}\cdot d\theta_{\mu,m}\cdot dP_{\mu,n}}\\
\tilde{F}_{lmn} &= \left(POT\cdot T\cdot\int_{lmn}\overline{F}\left(E_{\nu,l}\right)\cdot dE_{\nu,l}\cdot d\theta_{\mu,m}\cdot dP_{\mu,n}\right)\\
\tilde{\Delta}_{ijklmn} &= \frac{POT\cdot T\cdot\int_{lmn}F\left(E_{\nu,l}\right)\cdot\frac{d\sigma(E_{\nu,l},\theta_{\mu,m},P_{\mu,n})}{d\theta_{\theta,m}dP_{\mu,n}}\cdot D\cdot\epsilon\cdot dE_{\nu,l}\cdot d\theta_{\mu,m}\cdot dP_{\mu,n}}{POT\cdot T\cdot\int_{lmn}\overline{F}\left(E_{\nu,l}\right)\cdot\frac{d\sigma(E_{\nu,l},\theta_{\mu,m},P_{\mu,n})}{d\theta_{\mu,m}dP_{\mu,n}}\cdot dE_{\nu,l}\cdot d\theta_{\mu,m}\cdot dP_{\mu,n}}\\
s_{ijklmn} &= \tilde{\Delta}_{ijklmn}\cdot\tilde{F}_{lmn}\cdot s_{lmn}
\end{aligned} \tag{10.28}$$

By enumerating the reconstructed bins in $\{i,j,k\}$ (now using the index i) as well as the truth bins in $\{l,m,n\}$ (now using the index j), the master equation relating measurement to truth

bins can be constructed:

$$m_i - b_i = \sum_j \tilde{\Delta}_{ij} \cdot \tilde{F}_j \cdot s_j = \sum_j r_{ij} \cdot s_j \tag{10.29}$$

With this relation established, it is possible to unfold to the triple-differential result by first computing the flux constant and estimating the mapping from truth to reconstructed bins in simulation, and then performing the Wiener SVD unfolding as outlined earlier.

11 Cross Section Measurements

11.1 Single-Differential Cross Section Measurements

The triple-differential measurement central to this dissertation builds on the work done to extract single-differential measurements over neutrino energy, muon energy, and transfer energy [167]. Those results are discusses here, before proceeding to the triple-differential measurement. The single-differential analysis was performed on only 5.3×10^{20} POT of data (instead of the full 6.4×10^{20} POT available in runs 1-3) to comply with the MicroBooNE blindness policy in place at the time. However, statistical fluctuations are still small because of the relatively small number of bins in 1D compared to 3D.

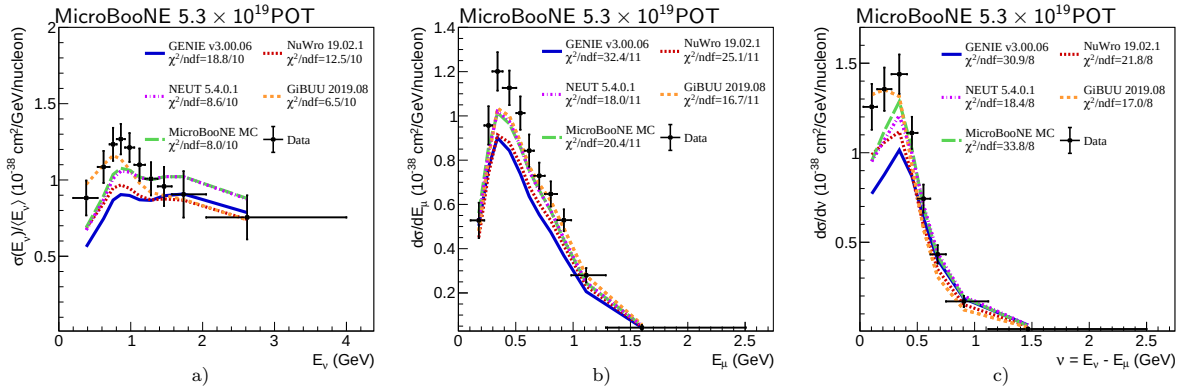

Figure 95: a) The extracted ν_μCC inclusive scattering cross section per nucleon divided by the bin-center neutrino energy, as a function of neutrino energy. b) The measured ν_μCC differential cross section per nucleon as a function of muon energy $d\sigma/dE_\mu$. c) The measured ν_μCC differential cross section per nucleon as a function of energy transfer $d\sigma/d\nu$. Various model predictions are compared to all three measurements (see text for details). Source [167]

In addition to the extracted cross section measurement and MicroBooNE model prediction, comparisons are shown against the predictions of Genie v3.0.6 (untuned), NuWro 19.02.01, NEUT 5.4.01, and GiBUU 2019.08. In each case, the additional smearing matrix A_C is applied to the model prediction for an unbiased comparison, with χ^2/ndf calculations shown alongside each model prediction in Fig. 95. The covariance matrices corresponding to each unfolded measurement are shown in Fig. 96. For each model comparison, the CV

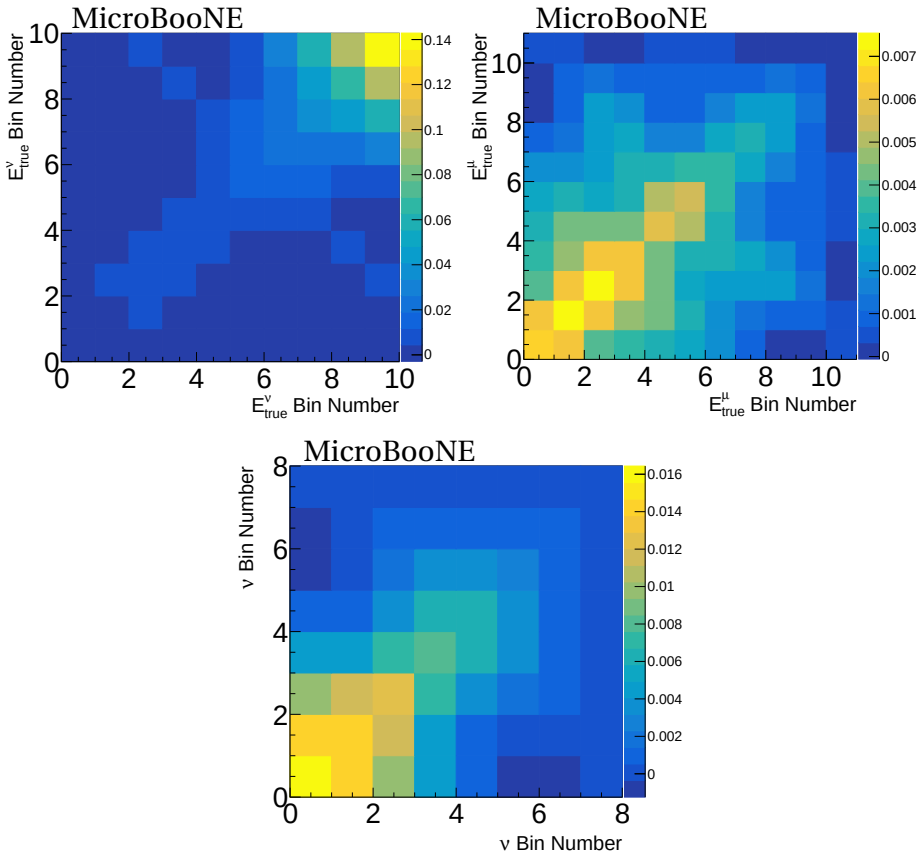

Figure 96: Total covariance matrices in different energy bins for (a) total cross section per nucleon as a function of neutrino energy, (b) differential cross section per nucleon as a function of muon energy, and (c) differential cross section per nucleon as a function of energy transfer. Source: [167]

prediction without uncertainties is used, as it is difficult to properly incorporate the uncertainties of each model. Model predictions for Genie v3 (untuned) and NuWro are dis-favored across the three measurements, while GiBUU consistently gives the best prediction. The differences in model prediction are most clear in Fig. 95c over ν, especially at low transfer energy where GiBUU performs particularly well.

11.2 Triple-Differential Cross Section Measurement

The full triple-differential cross section is shown in Fig. 97. Each angular slice is presented in a separate panel, arranged from backward to forward facing. Within each panel, the four E_ν

Table 5: Comparisons between various models and the unfolded triple-differential measurement, considered over the entire 138 analysis bins and within each E_ν slice, specified by the range in GeV.

Model Name	Total χ^2/ndf	[0.2, 0.705] χ^2/ndf	[0.705, 1.05] χ^2/ndf	[1.05, 1.57] χ^2/ndf	[1.57, 4.0] χ^2/ndf
GENIE v2.12.10	741.1/138	71.4/28	64.4/35	64.3/42	35.6/33
GENIE v3.0.6 (uBooNE tune)	326.1/138	85.0/28	77.8/35	44.6/42	31.9/33
GENIE v3.0.6 (untuned)	322.2/138	94.1/28	84.8/35	52.2/42	37.3/33
GiBUU 2021	269.9/138	33.8/28	54.8/35	52.6/42	31.0/33
NEUT 5.4.0.1	243.3/138	58.5/28	59.9/35	33.1/42	38.2/33
NuWro 19.02.01	212.1/138	54.8/28	67.3/35	40.9/42	29.6/33

slices are plotted on top of each other, given an arbitrary offset Δ that varies between slices to help visually separate each slice. Within each panel the measurement is shown over the P_μ bin centers within that slice. The triple-differential result provides a highly detailed picture of the physics involved in ν_μCC interactions, including the energy and angular dependence of the peak in cross section as a function of muon momentum. Additionally, the unfolding to E_ν helps separate interaction processes: NuWro predicts a reduction in the quasi-elastic fraction from $\sim 75\%$ in the lowest energy bin to $\sim 55\%$ in the highest energy bin.

Although the data is unfolded with the MicroBooNE model prediction, it is shown here in comparison to the NuWro 19.02.01 prediction as it has the best agreement with the data. The quality of fit for each model comparison is shown in Table 5, including comparisons with GENIE v2.12.10, GENIE 3.0.6 G18_10a_02_11a [145, 54], NuWro 19.02.01 [161], NEUT 5.4.0.1 [160], and GiBUU 2021 [170]. The differences in physics modeling assumptions used between the event generators are described in Ref. [171].

The triple-differential measurement provides data covering a large phase space, causing it to be in tension with all model CV predictions considered. The best overall model agreement is found with NuWro, followed by NEUT and GiBUU. The hierarchy of model agreement largely matches what was found in the single-differential results, with the notable exception that NuWro perform much better over this higher-dimensional phase space. To help visualize

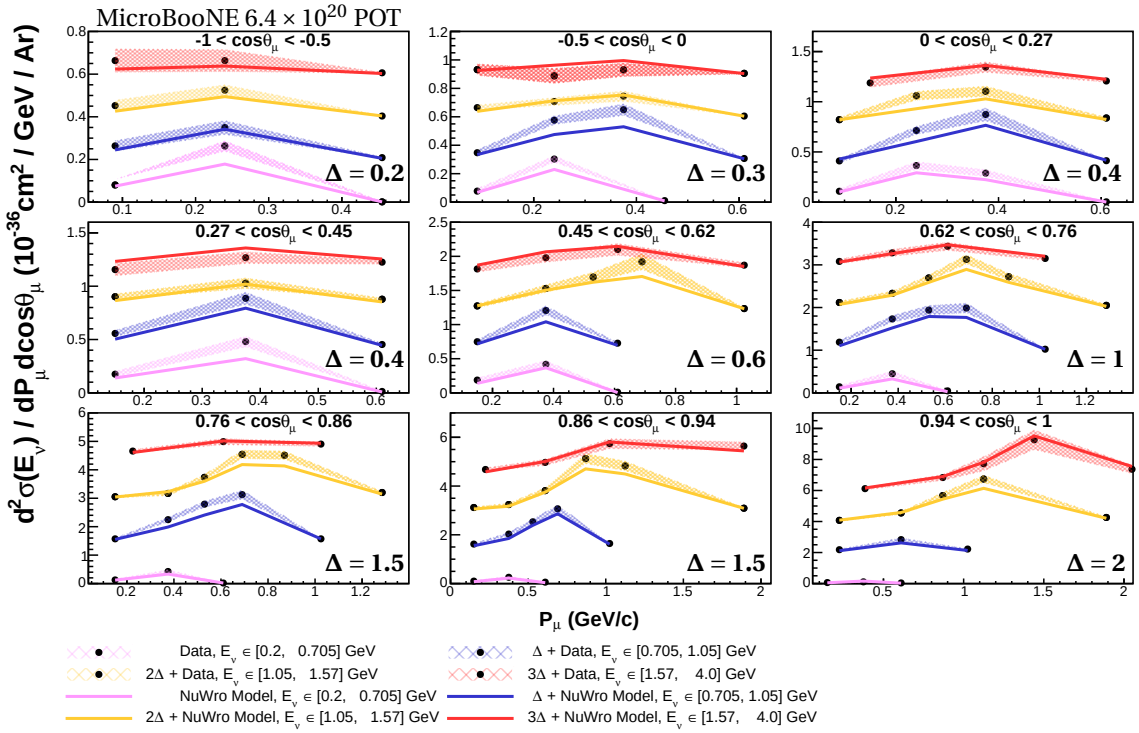

Figure 97: Unfolded differential cross section measurement and NuWro prediction, chosen for having the lowest χ^2, are shown over P_μ within each angle slice. Measurements from each E_ν slice are overlaid and offset to visually separate them. The magnitude of the offset Δ, given in the same units of 10^{-36}cm^2/GeV/Ar, is listed in the bottom right of each plot.

the data/MC comparisons, the muon momentum dimension is integrated over and each bin is normalized by the average neutrino energy in the slice, producing the double-differential result shown in Fig. 98. GiBUU provides the best description of the data below 1 GeV, but is outperformed at higher energies, particularly by NuWro. The high energy region, particularly in combination with low Q^2 and forward lepton angles, is notable for an increase in Δ-resonance interactions combined with differences in the modeling of pion production [172] that may play a role in the overall model performances.

In the future, this analysis can be enhanced through the inclusion of additional data, both from the BNB used in this analysis already, as well as the NuMI beamline that MicroBooNE is exposed to. Since the NuMI beam is off-axis and at a higher energy, the flux predictions between these two beamlines are only slightly correlated, meaning that a combined BNB+NuMI

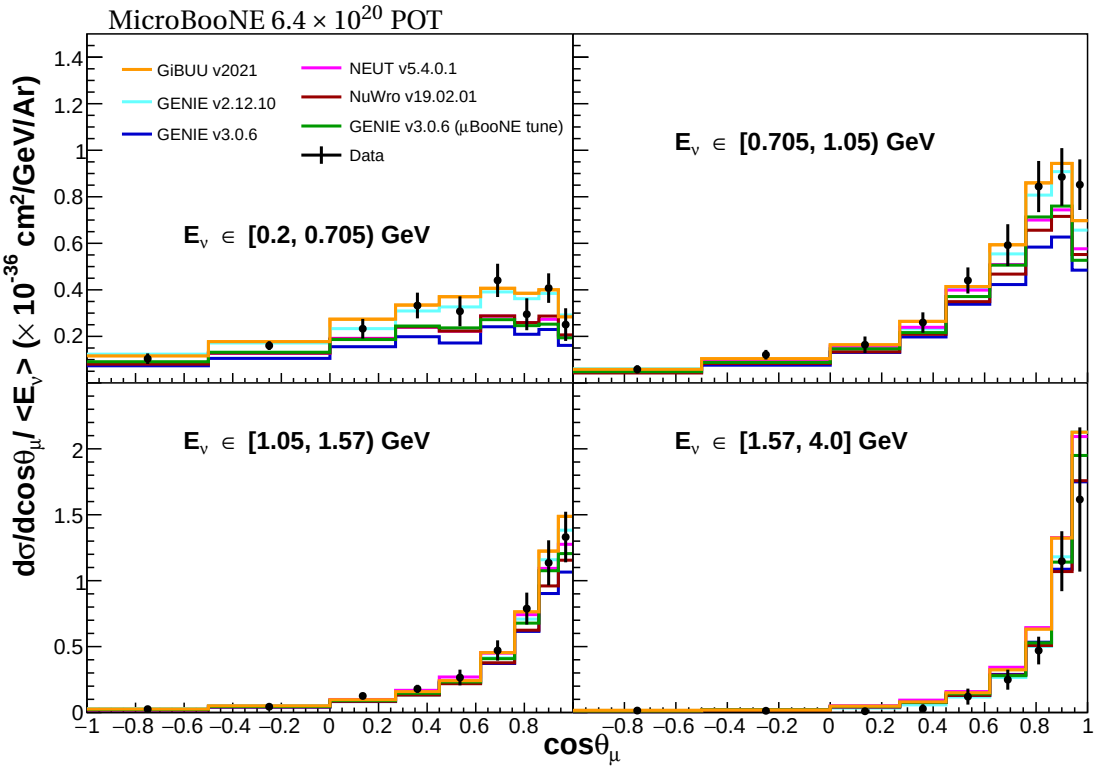

Figure 98: Unfolded differential cross section over $cos(\theta_\mu)$ after integrating over P_μ and normalizing by the average $\langle E_\nu \rangle$ in each E_ν bin.

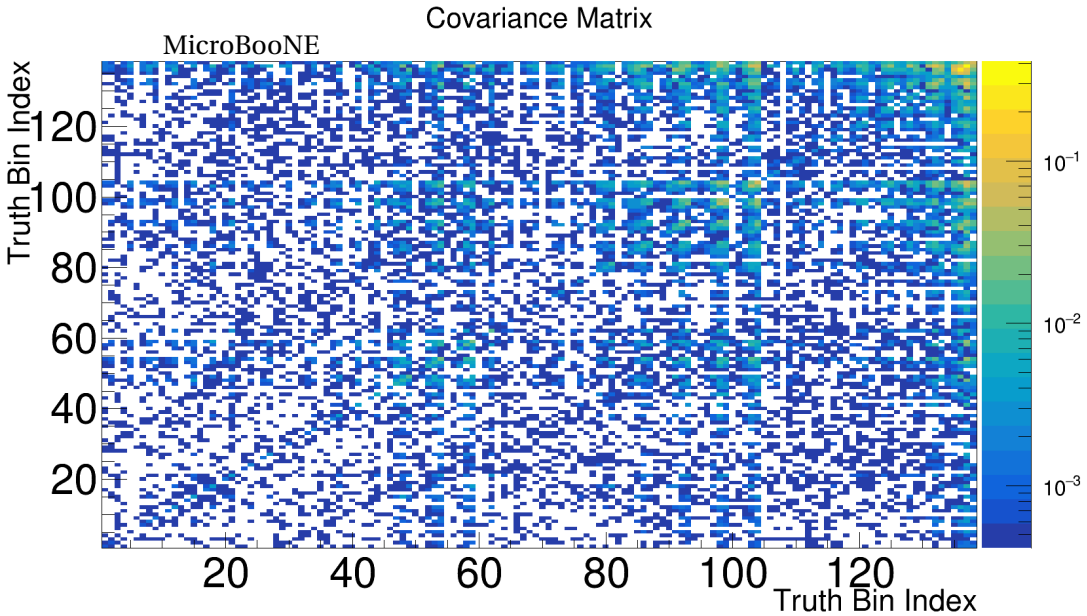

Figure 99: Covariance matrix showing total statistical plus systematic uncertainties over the 138 measurement bins, in units of 10^{-36} cm^2/Ar.

measurement will not only reduce the statistical uncertainties but also the flux uncertainties in the extracted measurement. In parallel with these advances, work is in-progress to produce multi-differential cross section measurements over exclusive channels that comprise the overall inclusive ν_μCC channel.

11.3 Conclusion

MicroBooNE has served as a pioneer in the field of LArTPC-based neutrino physics experiments, fulfilling primary goals of investigating the MiniBooNE LEE, providing research and development on large-scale LArTPC detectors, and producing cross section measurements. Within MicroBooNE, Wire-Cell has developed a reconstruction chain that reconstructs a fully 3D charge distribution and produces produces highly accurate data products on particle identities, kinematics, and overall organization within a neutrino interaction. Notable steps within this reconstruction include signal processing with 2D deconvolution, tomographic imaging for 3D charge reconstruction, global charge-light matching of TPC activity, graph-based trajectory fitting and particle flow, and neutrino vertex identification incorporating a deep neural network. Using this reconstruction, Wire-Cell developed a high-efficiency, high-purity inclusive ν_μCC selection that achieves 92% purity while maintaining 68% efficiency. This remarkable selection performance as well as reconstruction quality enables the use of a new data-driven validation procedure that leverages correlated model predictions across final state kinematics to increase sensitivity to mis-modeling, including over invisible TPC activity. Finally, Wiener SVD unfolding was used to produce a nominal-flux averaged differential inclusive ν_μ CC cross section $\frac{d^2\sigma(E_\nu)}{d\cos(\theta_\mu)dP_\mu}$. This measurement uses an exposure of 6.4×10^{20} POT of data from the BNB, and contains a wealth of info capable of assisting model development. In particular, the validated reconstruction of neutrino energy is especially useful for high-precision neutrino oscillation searches, such as at the upcoming DUNE experiment.

References

[1] C. Cowan, F. Reines, F. Harrison, F. Kruse, and A. McGuire. Detection of the Free Neutrino: a Confirmation. Science, 124:103–104, Jul 1956.

[2] Julian Schwinger. A theory of the fundamental interactions. Annals of Physics, 2(5):407–434, 1957.

[3] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos. Phys. Rev., 109:1015–1017, Feb 1958.

[4] W. N. Cottingham and D. A. Greenwood. An Introduction to the Standard Model of Particle Physics. Cambridge University Press, 2 edition, 2007.

[5] M. C. Gonzalez-Garcia and Yosef Nir. Neutrino masses and mixing: evidence and implications. Rev. Mod. Phys., 75:345–402, Mar 2003.

[6] Goran Senjanović and Vladimir Tello. Parity and the origin of neutrino mass. International Journal of Modern Physics A, 35(09):2050053, mar 2020.

[7] Hitoshi Murayama. Origin of neutrino mass. Progress in Particle and Nuclear Physics, 57(1):3–21, 2006. International Workshop of Nuclear Physics 27th course.

[8] Rabindra Mohapatra and Goran Senjanović. Neutrino mass and spontaneous parity nonconservation. Physical Review Letters - PHYS REV LETT, 44:912–915, 04 1980.

[9] Tsutomu Yanagida. Horizontal symmetry and mass of the t quark. Phys. Rev. D, 20:2986–2988, Dec 1979.

[10] J. Schechter and J. W. F. Valle. Neutrino masses in su(2) $\otimes$ u(1) theories. Phys. Rev. D, 22:2227–2235, Nov 1980.

[11] Howard Georgi and S. L. Glashow. Unity of all elementary-particle forces. Phys. Rev. Lett., 32:438–441, Feb 1974.

[12] Ettore Majorana and Luciano Maiani. A symmetric theory of electrons and positrons, pages 201–233. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

[13] J. Schechter and J. W. F. Valle. Neutrinoless double-β decay in su(2)×u(1) theories. Phys. Rev. D, 25:2951–2954, Jun 1982.

[14] D. Q. Adams et al. Improved limit on neutrinoless double-beta decay in ^{130}Te with cuore. Phys. Rev. Lett., 124:122501, Mar 2020.

[15] Werner Tornow. Search for Neutrinoless Double-Beta Decay, 2014.

[16] Makoto Kobayashi and Toshihide Maskawa. CP-Violation in the Renormalizable Theory of Weak Interaction. Progress of Theoretical Physics, 49(2):652–657, 02 1973.

[17] Particle Data Group. Review of Particle Physics. Progress of Theoretical and Experimental Physics, 2020(8), 08 2020. 083C01.

[18] Bruno Pontecorvo. Mesonium and antimesonium. Journal of Experimental and Theoretical Physics, 33:429, 1957.

[19] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the Unified Model of Elementary Particles. Progress of Theoretical Physics, 28(5):870–880, 11 1962.

[20] J W F Valle. Neutrino physics overview. Journal of Physics: Conference Series, 53(1):473, nov 2006.

[21] X. Qian and P. Vogel. Neutrino mass hierarchy. Progress in Particle and Nuclear Physics, 83:1–30, 2015.

[22] The T2K Collaboration. Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations. Nature, 580(7803):339–344, Apr 2020.

[23] C. R. Das, João Pulido, Jukka Maalampi, and Sampsa Vihonen. Determination of the θ_{23} octant in long baseline neutrino experiments within and beyond the standard model. Phys. Rev. D, 97:035023, Feb 2018.

[24] M. Aker et al. Improved upper limit on the neutrino mass from a direct kinematic method by katrin. Phys. Rev. Lett., 123:221802, Nov 2019.

[25] G. L. Fogli, E. Lisi, A. Marrone, D. Montanino, A. Palazzo, and A. M. Rotunno. Global analysis of neutrino masses, mixings, and phases: Entering the era of leptonic cp violation searches. Phys. Rev. D, 86:013012, Jul 2012.

[26] P.S. Barbeau. Coherent neutrino scattering. Nuclear and Particle Physics Proceedings, 265-266:117–119, 2015. Proceedings of the Neutrino Oscillation Workshop.

[27] Jan T Sobczyk. Quasi-elastic neutrino scattering - an overview. AIP Conference Proceedings, 1405(1), 11 2011.

[28] A. M. Ankowski and et al. Electron scattering and neutrino physics. TBD, 3 2022.

[29] Ger van Middelkoop. Foreword. In S. Bentvelsen, P. de Jong, J. Koch, and E. Laenen, editors, Proceedings of the 31st International Conference on High Energy Physics Ichep 2002. JAI, Amsterdam, 2003.

[30] L. David Roper. Evidence for a P_{11} pion-nucleon resonance at 556 mev. Phys. Rev. Lett., 12:340–342, Mar 1964.

[31] J. A. Formaggio and G. P. Zeller. From ev to eev: Neutrino cross sections across energy scales. Rev. Mod. Phys., 84:1307–1341, Sep 2012.

[32] P. Vilain et al. Coherent single charged pion production by neutrinos. Physics Letters B, 313(1):267–275, 1993.

[33] Zelimir Djurcic. Backgrounds in neutrino appearance signal at miniboone. AIP Conference Proceedings, 842(1), 7 2006.

[34] Janet M. Conrad, Michael H. Shaevitz, and Tim Bolton. Precision measurements with high-energy neutrino beams. Rev. Mod. Phys., 70:1341–1392, Oct 1998.

[35] A. Kayis-Topaksu, Pierre Vilain, and Gaston Wilquet. Leading order analysis of neutrino induced dimuon events in the chorus experiment. Nuclear Physics, 798:1–16, 2008.

[36] D. Mason et al. Measurement of the nucleon strange-antistrange asymmetry at next-to-leading order in qcd from nutev dimuon data. Phys. Rev. Lett., 99:192001, Nov 2007.

[37] Q. Wu et al. A precise measurement of the muon neutrino–nucleon inclusive charged current cross section off an isoscalar target in the energy range $2.5 < E_\nu < 40$ GeV by NOMAD. Physics Letters B, 660(1):19–25, 2008.

[38] P. Adamson et al. Neutrino and antineutrino inclusive charged-current cross section measurements with the minos near detector. Phys. Rev. D, 81:072002, Apr 2010.

[39] J. Mousseau et al. Measurement of partonic nuclear effects in deep-inelastic neutrino scattering using minerva. Phys. Rev. D, 93:071101, Apr 2016.

[40] Omar Benhar, Adelchi Fabrocini, and Stefano Fantoni. The nucleon spectral function in nuclear matter. Nuclear Physics A, 505(2):267–299, 1989.

[41] H. Hassanabadi, A. Armat, and L. Naderi. Relativistic fermi-gas model for nucleus. Foundations of Physics, 44(11):1188–1194, Nov 2014.

[42] J. Nieves, J. E. Amaro, and M. Valverde. Inclusive quasielastic charged-current neutrino-nucleus reactions. Phys. Rev. C, 70:055503, Nov 2004.

[43] Omar Benhar, Nicola Farina, Hiroki Nakamura, Makoto Sakuda, and Ryoichi Seki. Electron- and neutrino-nucleus scattering in the impulse approximation regime. Phys. Rev. D, 72:053005, Sep 2005.

[44] Ranjeet Dalal and I. J. Douglas MacGregor. Nucleon-nucleon correlations inside atomic nuclei: synergies, observations and theoretical models, 2022.

[45] E. Piasetzky, M. Sargsian, L. Frankfurt, M. Strikman, and J. W. Watson. Evidence for strong dominance of proton-neutron correlations in nuclei. Phys. Rev. Lett., 97:162504, Oct 2006.

[46] I. Ruiz Simo, J.E. Amaro, M.B. Barbaro, A. De Pace, J.A. Caballero, G.D. Megias, and T.W. Donnelly. Emission of neutron–proton and proton–proton pairs in neutrino scattering. Physics Letters B, 762:124–130, 2016.

[47] J. Nieves, J. E. Amaro, and M. Valverde. Inclusive quasielastic charged-current neutrino-nucleus reactions. Phys. Rev. C, 70:055503, Nov 2004.

[48] K. Abe et al. First t2k measurement of transverse kinematic imbalance in the muon-neutrino charged-current single-π^+ production channel containing at least one proton. Phys. Rev. D, 103:112009, Jun 2021.

[49] Lars Bathe-Peters, Steven Gardiner, and Roxanne Guenette. Comparing generator predictions of transverse kinematic imbalance in neutrino-argon scattering, 2022.

[50] C. Andreopoulos, A. Bell, D. Bhattacharya, F. Cavanna, J. Dobson, S. Dytman, H. Gallagher, P. Guzowski, R. Hatcher, P. Kehayias, A. Meregaglia, D. Naples, G. Pearce, A. Rubbia, M. Whalley, and T. Yang. The genie neutrino monte carlo generator. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 614(1):87–104, 2010.

[51] J. Cugnon. Intranuclear cascade model. a review. Nuclear Physics A, 387(1):191–203, 1982.

[52] M Kim, H Bhang, Jeonghun Kim, Young-Seog Kim, Hyeonseo Park, and J Chang. A monte-carlo intranuclear cascade calculation for the propagation of energetic nucleons in the nucleus. Journal of the Korean Physical Society, 46, 04 2005.

[53] S. Dytman, Y. Hayato, R. Raboanary, J. T. Sobczyk, J. Tena-Vidal, and N. Vololoniaina. Comparison of validation methods of simulations for final state interactions in hadron production experiments. Phys. Rev. D, 104:053006, Sep 2021.

[54] Luis Alvarez-Ruso et al. Recent highlights from genie v3. The European Physical Journal Special Topics, 230(24):4449–4467, Dec 2021.

[55] Tomasz Golan, Cezary Juszczak, and Jan T. Sobczyk. Effects of final-state interactions in neutrino-nucleus interactions. Phys. Rev. C, 86:015505, Jul 2012.

[56] Bruce T. Cleveland, Timothy Daily, Jr. Raymond Davis, James R. Distel, Kenneth Lande, C. K. Lee, Paul S. Wildenhain, and Jack Ullman. Measurement of the solar electron neutrino flux with the homestake chlorine detector. The Astrophysical Journal, 496(1):505, mar 1998.

[57] Y. Fukuda et al. Study of the atmospheric neutrino flux in the multi-gev energy range. Physics Letters B, 436(1):33–41, 1998.

[58] T. Kajita, E. Kearns, and M. Shiozawa. Establishing atmospheric neutrino oscillations with super-kamiokande. Nuclear Physics B, 908:14–29, 2016. Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015.

[59] Q. R. Ahmad et al. Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8b solar neutrinos at the sudbury neutrino observatory. Phys. Rev. Lett., 87:071301, Jul 2001.

[60] Srubabati Goswami. Solar neutrino experiments: An overview, 2003.

[61] S. P. Mikheyev and A. Yu. Smirnov. Resonant amplification of ν oscillations in matter and solar-neutrino spectroscopy. Il Nuovo Cimento C, 9(1):17–26, Jan 1986.

[62] L. Wolfenstein. Neutrino oscillations in matter. Phys. Rev. D, 17:2369–2374, May 1978.

[63] J. Seguinot and T. Ypsilantis. Photo-ionization and cherenkov ring imaging. Nuclear Instruments and Methods, 142(3):377–391, 1977.

[64] Yuichi Oyama. Results from kamiokande-ii experiment. Nuclear Physics B - Proceedings Supplements, 13:362–367, 1990.

[65] Katharina Lodders. Solar system abundances and condensation temperatures of the elements. The Astrophysical Journal, 591(2):1220, jul 2003.

[66] R. L. William. Techniques for Nuclear and Particle Physics Experiments. Springer Berlin, 2 edition, 1994.

[67] A. Gando et al. Constraints on θ_{13} from a three-flavor oscillation analysis of reactor antineutrinos at kamland. Phys. Rev. D, 83:052002, Mar 2011.

[68] F. P. An et al. Observation of electron-antineutrino disappearance at daya bay. Phys. Rev. Lett., 108:171803, Apr 2012.

[69] Y. Abe et al. Reactor $\overline{\nu}_e$ disappearance in the double chooz experiment. Phys. Rev. D, 86:052008, Sep 2012.

[70] J. K. Ahn et al. Observation of reactor electron antineutrinos disappearance in the reno experiment. Phys. Rev. Lett., 108:191802, May 2012.

[71] R. P. Litchfield. (direct) measurement of theta-13, 2012.

[72] Daniel A. Dwyer. The neutrino mixing angle θ13: Reactor and accelerator experiments. Physics of the Dark Universe, 4:31–35, 2014. DARK TAUP2013.

[73] R. Schwienhorst et al. A new upper limit for the tau-neutrino magnetic moment. Physics Letters B, 513(1):23–29, 2001.

[74] André de Gouvêa, Kevin J. Kelly, G. V. Stenico, and Pedro Pasquini. Physics with beam tau-neutrino appearance at DUNE. Physical Review D, 100(1), jul 2019.

[75] K. Abe et al. Indication of electron neutrino appearance from an accelerator-produced off-axis muon neutrino beam. Phys. Rev. Lett., 107:041801, Jul 2011.

[76] K. Abe et al. Precise measurement of the neutrino mixing parameter θ_{23} from muon neutrino disappearance in an off-axis beam. Phys. Rev. Lett., 112:181801, May 2014.

[77] Melissa George. Electron neutrinos at t2k, 2010.

[78] M. C. Gonzalez-Garcia, Michele Maltoni, Jordi Salvado, and Thomas Schwetz. Global fit to three neutrino mixing: critical look at present precision. Journal of High Energy Physics, 2012(12):123, Dec 2012.

[79] N. Agafonova et al. Observation of a first ν_τ candidate event in the opera experiment in the cngs beam. Physics Letters B, 691(3):138–145, 2010.

[80] Alessandro Strumia. Interpreting the lsnd anomaly: sterile neutrinos or cpt-violation or...? Physics Letters B, 539(1):91–101, 2002.

[81] Scott Dodelson, Alessandro Melchiorri, and An že Slosar. Is cosmology compatible with sterile neutrinos? Phys. Rev. Lett., 97:041301, Jul 2006.

[82] A. A. Aguilar-Arevalo et al. Significant excess of electronlike events in the miniboone short-baseline neutrino experiment. Phys. Rev. Lett., 121:221801, Nov 2018.

[83] P. Abratenko et al. Search for an excess of electron neutrino interactions in microboone using multiple final-state topologies. Phys. Rev. Lett., 128:241801, Jun 2022.

[84] Peter Galison. Image and logic: a material culture of microphysics. Chapter 3, Nuclear Emulsions: The Anxiety of the Experimenter. University of Chicago Press, 1997.

[85] Donald A. Glaser. Some effects of ionizing radiation on the formation of bubbles in liquids. Phys. Rev., 87:665–665, Aug 1952.

[86] W.J. Willis and V. Radeka. Liquid-argon ionization chambers as total-absorption detectors. Nuclear Instruments and Methods, 120(2):221–236, 1974.

[87] 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, Mar 2022.

[88] I. Stancu. Technical design report for the 8 gev beam, 5 2001.

[89] A. A. Aguilar-Arevalo et al. Neutrino flux prediction at miniboone. Phys. Rev. D, 79:072002, Apr 2009.

[90] None None. Booster neutrino flux prediction at microboone, 7 2018.

[91] MicroBooNE. A measurement of the attenuation of drifting electrons in the microboone lartpc, 8 2017.

[92] Ettore Segreto. Properties of liquid argon scintillation light emission. Phys. Rev. D, 103:043001, Feb 2021.

[93] R Acciarri et al. A study of electron recombination using highly ionizing particles in the argoneut liquid argon tpc. Journal of Instrumentation, 8(08):P08005, aug 2013.

[94] Yichen Li, Thomas Tsang, Craig Thorn, Xin Qian, Milind Diwan, Jyoti Joshi, Steve Kettell, William Morse, Triveni Rao, James Stewart, Wei Tang, and Brett Viren. Measurement of longitudinal electron diffusion in liquid argon. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 816:160–170, 2016.

[95] P. Abratenko et al. Cosmic ray background rejection with wire-cell lartpc event reconstruction in the microboone detector. Phys. Rev. Applied, 15:064071, Jun 2021.

[96] R. Acciarri et al. Design and construction of the microboone detector. Journal of Instrumentation, 12(02):P02017, feb 2017.

[97] B.K. Lubsandorzhiev. On the history of photomultiplier tube invention. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 567(1):236–238, 2006. Proceedings of the 4th International Conference on New Developments in Photodetection.

[98] Rob Veenhof. Garfield, recent developments. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 419(2):726–730, 1998.

[99] C. Adams et al. Ionization electron signal processing in single phase lartpcs. part i. algorithm description and quantitative evaluation with microboone simulation. Journal of Instrumentation, 13(07):P07006, jul 2018.

[100] G. Cavalleri, E. Gatti, G. Fabri, and V. Svelto. Extension of ramo's theorem as applied to induced charge in semiconductor detectors. Nuclear Instruments and Methods, 92(1):137–140, 1971.

[101] C. Adams et al. Ionization electron signal processing in single phase lartpcs. part ii. data/simulation comparison and performance in microboone. Journal of Instrumentation, 13(07):P07007, jul 2018.

[102] B. Baller. Liquid argon tpc signal formation, signal processing and reconstruction techniques. Journal of Instrumentation, 12(07):P07010, jul 2017.

[103] R. Acciarri et al. Noise characterization and filtering in the microboone liquid argon tpc. Journal of Instrumentation, 12(08):P08003, aug 2017.

[104] Norbert Wiener. Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications. The MIT Press, 1949.

[105] Willa W Chen, Clifford M Hurvich, and Yi Lu. On the correlation matrix of the discrete fourier transform and the fast solution of large toeplitz systems for long-memory time series. Journal of the American Statistical Association, 101(474):812–822, 2006.

[106] 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. Journal of Instrumentation, 16(06):P06043, jun 2021.

[107] X. Qian, C. Zhang, B. Viren, and M. Diwan. Three-dimensional imaging for large lartpcs. Journal of Instrumentation, 13(05):P05032, may 2018.

[108] Peter A. Burrough, Rachael McDonnell, Rachael A. McDonnell, and Christopher D. Lloyd. Principles of Geographical Information Systems. Oxford University Press, 2015.

[109] Michael Lustig, David Donoho, and John M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. Magnetic Resonance in Medicine, 58, oct 2007.

[110] Emmanuel Candes, Justin Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements, 2005.

[111] Hong Cheng. The Fundamentals of Compressed Sensing, pages 21–53. Springer London, London, 2015.

[112] J. L. Blanco and P. K. Rai. anoflann: a c++ header-only fork of flann, a library for nearest neighbor (nn) wih kd-trees, 2014.

[113] P V.C. Hough. Method and means for recognizing complex patterns, 12 1962.

[114] I. T. Jolliffe. Principal Component Analysis. Springer New York, 2 edition, 2002.

[115] akuukka. quickhull, 2017.

[116] S. L. Lauritzen. Thiele: Pioneer in Statistics. Oxford University Press, 2002.

[117] E. W. Dijkstra. A note on two problems in connexion with graphs. Numerische Mathematik, 1(1):269–271, Dec 1959.

[118] M. Heideman, D. Johnson, and C. Burrus. Gauss and the history of the fast fourier transform. IEEE ASSP Magazine, 1(4):14–21, 1984.

[119] M. A. Stephens. Edf statistics for goodness of fit and some comparisons. Journal of the American Statistical Association, 69(347):730–737, 1974.

[120] Benjamin J. P. Jones. Sterile neutrinos in cold climates, 2015.

[121] M. Antonello et al. Precise 3d track reconstruction algorithm for the icarus t600 liquid argon time projection chamber detector. Advances in High Energy Physics, 2013:260820, Mar 2013.

[122] Marshall W Bern and Ronald L. Graham. The shortest-network problem. Scientific American, 260:84–89, Jan 1989.

[123] Piotr Wygocki. Steiner tree greedy algorithm, 2019.

[124] Biconjugate gradient stablized method (bicgstab), 2019.

[125] P. Abratenko et al. Wire-cell 3d pattern recognition techniques for neutrino event reconstruction in large lartpcs: algorithm description and quantitative evaluation with microboone simulation. Journal of Instrumentation, 17(01):P01037, jan 2022.

[126] P. Abratenko et al. Determination of muon momentum in the microboone lartpc using an improved model of multiple coulomb scattering. Journal of Instrumentation, 12(10):P10010, oct 2017.

[127] Jason Brownlee. A gentle introduction to object recognition with deep learning, 2019.

[128] P. Abratenko et al. Semantic segmentation with a sparse convolutional neural network for event reconstruction in microboone. Phys. Rev. D, 103:052012, Mar 2021.

[129] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9224–9232, 2018.

[130] Benjamin Graham and Laurens van der Maaten. Submanifold sparse convolutional networks, 2017.

[131] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, pages 234–241, Cham, 2015. Springer International Publishing.

[132] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields, 2016.

[133] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014.

[134] C. Adams et al. Reconstruction and measurement of o(100) mev energy electromagnetic activity from $\pi^0 \rightarrow \gamma\gamma$ decays in the microboone lartpc. Journal of Instrumentation, 15(02):P02007, feb 2020.

[135] PSTAR at NIST.

[136] P. Abratenko et al. High-performance generic neutrino detection in a lartpc near the earth's surface with the microboone detector, 2020.

[137] 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, 105:112005, Jun 2022.

[138] R. Acciarri et al. Michel electron reconstruction using cosmic-ray data from the microboone lartpc. Journal of Instrumentation, 12(09):P09014, sep 2017.

[139] Tianqi Chen and Carlos Guestrin. XGBoost. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, aug 2016.

[140] P. Abratenko et al. First Measurement of Inclusive Muon Neutrino Charged Current Differential Cross Sections on Argon at E_ν 0.8 GeV with the MicroBooNE Detector. Phys. Rev. Lett., 123:131801, 2019.

[141] Morris L. Eaton. Multivariate Statistics: A Vector Space Approach. John Wiley & Sons Inc, 1983.

[142] C. A. Argüelles, A. Schneider, and T. Yuan. A binned likelihood for stochastic models. Journal of High Energy Physics, 2019(6):30, Jun 2019.

[143] S. Agostinelli et al. Geant4—a simulation toolkit. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 506(3):250–303, 2003.

[144] G. Prior. The harp experiment: first physics results. Nuclear Physics A, 752:24–33, 2005. Proceedings of the 22nd International Nuclear Physics Conference (Part 2).

[145] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. Nucl. Instrum. Meth. A, 614:87–104, 2010.

[146] The MicroBooNE Collaboration. Neutrino interaction model and uncertainties for microboone analyses, 5 2020.

[147] K. 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, 93:112012, Jun 2016.

[148] P. Abratenko et al. New CC0π genie model tune for microboone. Phys. Rev. D, 105:072001, Apr 2022.

[149] 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.

[150] Byron P.Roe. Statistical errors in Monte Carlo estimates of systematic errors. Nuclear Instruments and Methods in Physics Research Section A, 570:159–164, 2006.

[151] P. Abratenko et al. Novel approach for evaluating detector-related uncertainties in a lartpc using microboone data. The European Physical Journal C, 82(5):454, May 2022.

[152] Chuong B. Do. Gaussian processes, 2007.

[153] Carl Edward Rasmussen and Christopher K. I. Williams. Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning). The MIT Press, 2005.

[154] Meghan Frate, Kyle Cranmer, Saarik Kalia, Alexander Vandenberg-Rodes, and Daniel Whiteson. Modeling smooth backgrounds and generic localized signals with gaussian processes, 2017.

[155] Quentin Baghi, Gilles Métris, Joël Bergé, Bruno Christophe, Pierre Touboul, and Manuel Rodrigues. Gaussian regression and power spectral density estimation with missing data: The microscope space mission as a case study. Phys. Rev. D, 93:122007, Jun 2016.

[156] Lingge Li, Nitish Nayak, Jianming Bian, and Pierre Baldi. Efficient neutrino oscillation parameter inference using Gaussian processes. Phys. Rev. D, 101:012001, Jan 2020.

[157] A. Bozswan et al. Unfolding with Gaussian Processes. arXiv:1811.01242, 2018.

[158] M. Frate et al. Title here. arXiv:1709.05681, 2017.

[159] Robert D. Cousins, Samuel J. May, and Yipeng Sun. Should unfolded histograms be used to test hypotheses?, 2016.

[160] Yoshinari Hayato. A neutrino interaction simulation program library NEUT. Acta Phys. Polon. B, 40:2477–2489, 2009.

[161] T. Golan, J. T. Sobczyk, and J. Zmuda. NuWro: the Wroclaw Monte Carlo Generator of Neutrino Interactions. Nucl. Phys. B Proc. Suppl., 229-232:499–499, 2012.

[162] Ulrich Mosel. Neutrino event generators: foundation, status and future. J. Phys. G, 46(11):113001, 2019.

[163] P. A. Rodrigues et al. Identification of nuclear effects in neutrino-carbon interactions at low three-momentum transfer. Phys. Rev. Lett., 116:071802, Feb 2016.

[164] M. V. Ascencio et al. Measurement of inclusive charged-current ν_μ scattering on hydrocarbon at $< E_\nu > \sim 6$ GeV with low three-momentum transfer. Phys. Rev. D, 106:032001, Aug 2022.

[165] T. Hauschild and M. Jentschel. Comparison of maximum likelihood estimation and chi-square statistics applied to counting experiments. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 457(1):384–401, 2001.

[166] 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. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 961:163677, 2020.

[167] P. Abratenko et al. First Measurement of Inclusive Muon Neutrino Charged Current Differential Cross Sections on Argon at E_ν ~0.8 GeV with the MicroBooNE Detector. Phys. Rev. Lett., 123(13):131801, 2019.

[168] Andre-Louis Cholesky. Note sur une méthode de résolution des équations normales provenant de l'application de la méthode des moindres carrés a un système d'équations linéaires en nombre inférieur a celui des inconnues. — application de la méthode a la résolution d'un système defini d'équations linéaires. Bulletin géodésique, 2(1):67–77, Apr 1924.

[169] W. Tang, X. Li, X. Qian, H. Wei, and C. Zhang. Data Unfolding with Wiener-SVD Method. JINST, 12(10):P10002, 2017.

[170] O. Buss, T. Gaitanos, K. Gallmeister, H. van Hees, M. Kaskulov, O. Lalakulich, A. B. Larionov, T. Leitner, J. Weil, and U. Mosel. Transport-theoretical Description of Nuclear Reactions. Phys. Rept., 512:1–124, 2012.

[171] M. Buizza Avanzini et al. Comparisons and challenges of modern neutrino-scattering experiments. Phys. Rev. D, 105(9):092004, 2022.

[172] P. Stowell et al. Tuning the genie pion production model with minerva data. Phys. Rev. D, 100:072005, Oct 2019.