Supplemental Material

The following is supplemental material expanding upon the primary text presented in "Search for an anomalous excess of charged-current quasi-elastic ν_e interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction." Each section below lists the section of the primary text that is most relevant.

I. FURTHER DISCUSSION OF THE BDT ENSEMBLE

This section provides supplementary information related to Secs. VI.A and VI.B of the primary text. We have performed a number of tests with the purpose of validating the performance of the BDT ensemble. We summarize the results of these tests in this section.

We first examined the signal selection power of a BDT ensemble trained on simulated events from a given Run period when used to infer on simulated events from a different Run period. This is useful for two reasons. First, it provides a method for testing the ensemble on events that were in neither the training nor the validation sample of the constituent BDTs. Testing on a sample independent of the validation sample is important because the BDT training was halted when the classification error on the validation set didn't improve after 50 training iterations. Second, it tests the impact of removing training events from the MC prediction. When using the BDT ensemble of a given Run period to evaluate the signal likelihood of events from that run period, one must remove events which appeared in the training sample of each constituent BDT. In this study we use the BDT ensemble trained on run period 2 to evaluate simulated ν_e events from Run period 3 and vice versa, so one does not need to remove training events. Figures 1 and 2 show the results of this Run period swap study. As one can see, the differences between the selection using the correct and incorrect Run period ensemble are small. The correct run period performs slightly better, as expected due to small differences in the detector status specific to each run period. However, the relatively small change in ν_e selection efficiency suggests that the BDT ensemble is able to perform well on simulated events that do not appear in either the training or validation sample of constituent BDTs.

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

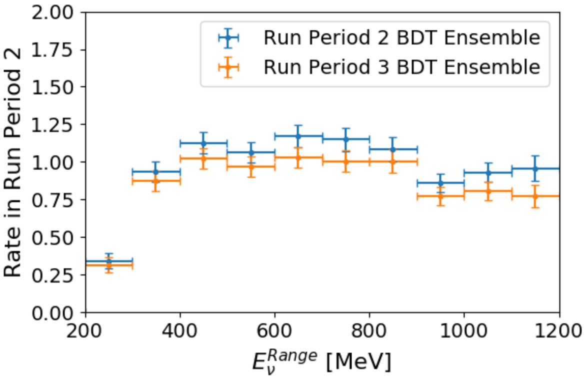

FIG. 1: The predicted ν_e event rate in Run period 2 using the Run period 2 ensemble (blue) and the Run period 3 ensemble (orange).

To this end, consider S_n to be the $1e1p$ BDT ensemble average score after removing n BDTs from the ensemble. For events in data, the BDT score is S_0 while for simulated events, the BDT score is S_n for some $n \in \{1, ..., 20\}$. Figs. 3 and 4 show the fractional difference $(S_n - S_0)/S_0$ as a function of the number of omitted BDTs n over signal-like (BDT score > 0.95) events in the simulation from Run period 2 and Run period 3, respectively. As in the previous study, in order to avoid bias from BDT training we use the BDT ensemble trained on Run period 2 to evaluate the signal likelihood of simulation events from Run period 3 and vice versa. The data points and error bars in Figs. 3 and 4 indicate the average and standard deviation of the fractional difference over the simulation sample, respectively. The red histograms show the actual distribution of omitted BDTs over the simulation from each Run period. One can see that the average BDT score does not exhibit significant bias upon removing BDTs from the calculation. Also, the fractional difference only becomes $> 1\%$ when removing $\gtrsim 18$ BDTs from the calculation. This only happens for $<< 1\%$ of simulated events, as expected.

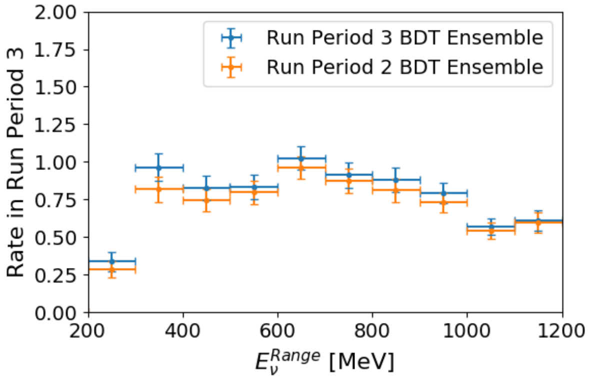

FIG. 2: The predicted ν_e event rate in Run period 3 using the Run period 3 ensemble (blue) and the Run period 2 ensemble (orange).

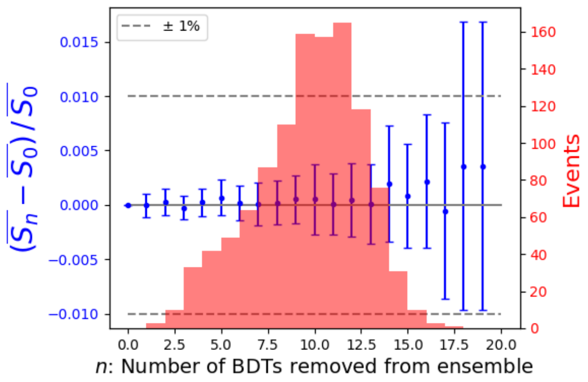

FIG. 4: The fractional difference in average BDT score $(S_n - S_0)/S_0$ as a function of the number of omitted BDTs n over the simulation from Run period 3. The data points and error bars indicate the average and standard deviation of the fractional difference over the simulation sample, respectively. The red histogram shows the actual distribution of the number of omitted BDTs over the Run period 3 simulation sample. Scores are calculated using the Run period 2 BDT ensemble.

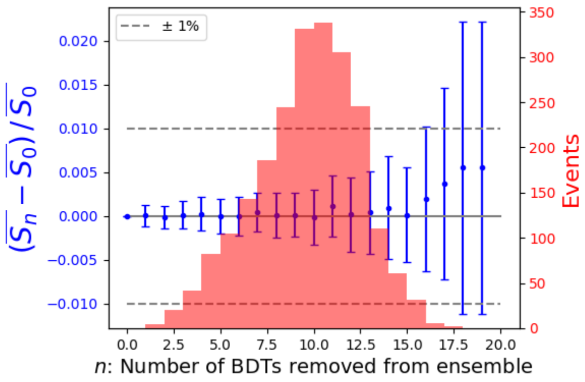

FIG. 3: The fractional difference in average BDT score $(S_n - S_0)/S_0$ as a function of the number of omitted BDTs n over the simulation from Run period 2. The data points and error bars indicate the average and standard deviation of the fractional difference over the simulation sample, respectively. The red histogram shows the actual distribution of the number of omitted BDTs over the Run period 2 simulation sample. Scores are calculated using the Run period 3 BDT ensemble.

II. FURTHER DISCUSSION ON BACKGROUNDS TO THE ANALYSIS

This section provides supplementary information related to Sec. VI.A.3 of the primary text. In this section we discuss the main background modes to the ν_e CCQE $1e1p$ analysis. These modes comprise the total background prediction which we fit to a Landau+linear function, as described in Sec. VI.A.3 of the primary text. The background modes include but are not limited to:

- ν_e charged-current interactions producing more than 1 proton in the final state, where the primary proton track exhibits CCQE-like kinematics.
- ν_μ interactions with a π^0 in the final state, where one of the photon showers from the π^0 decay is mistaken for an electron.
- Events with a vertex $> 5\,\mathrm{cm}$ away from a true neutrino vertex, which can happen when a photon shower from a π^0 decay falls near a cosmogenic muon, reconstructing a good one-track-one-shower event, or $\mu^\pm \to e^\pm$ decays, in which the end of the muon track is mistaken for a proton. These events are referred to as “off-vertex” events.

We now give more details on the Landau+linear probability distribution function (PDF) used to model the background contribution. The predicted background spectrum in each reconstructed neutrino energy bin $f(E_i)$ is generated by integrating the Landau+linear PDF $p(E)$ within that bin. Specifically,

$$\begin{aligned} p(E) &= \exp\left[\frac{-(E' + e^{-E'})}{2}\right] + aE, \\ E' &= (E-\mu)/\sigma, \\ f(E_i) &= \int_{E_i-\delta E_i}^{E_i+\delta E_i} p(E)dE, \end{aligned} \tag{1}$$

where μ and σ are the center and width of the Moyal approximation of the Landau function [1], a is the linear slope parameter, E_i is the center of the i’th energy bin, and δE_i is half of the bin width.

As described in Sec. VI.A.3, we estimate the overall background rate by fitting the BDT score distribution of simulated background events to a linear PDF $\tilde{p}(x)$. The integral of this PDF gives total expected background rate $\tilde{f}(x)$ at a BDT score cutoff at x:

$$\tilde{p}(x) = mx + b \qquad \tilde{f}(x) = \int_x^1 \tilde{p}(x)dx \tag{2}$$

where m and b are the slope and bias parameter for the BDT score distribution, respectively.

This procedure is intended to harness higher-statistics information on the background shape from a looser requirement on the $1e1p$ BDT score. This gives the background prediction shown in Figs. 4 and 16 of the primary text. However, the fit does not retain information on the predicted rate of specific background channels, such as the ν_μ π^0 and off-vertex events mentioned above.

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

TABLE I: Breakdown of events in the “background” category of Fig. 4 of the primary text over the range $200 < E_\nu < 1200\,\mathrm{MeV}$. The events are partitioned both by the interaction channel and the event topology. The prediction here comes directly from the simulation and does not incorporate the Landau+linear fit described in Sec. VI.A.3.

In Table I we give a breakdown of events in the “background” category of Fig. 4 of the primary text by their interaction channel and event topology. Note that these predicted rates come directly from the simulation rather than the Landau+linear fit, so the total predicted rate here (2.7) differs slightly from the total predicted rate using the background fit (3.2). However, this breakdown is illustrative of the relative predicted rates of different background channels. In Table I one can see that ν_μ resonant π^0 events are the most common background interaction channel, followed by off-vertex events. One can also see that events with a π^0 in the final state are the dominant background topology.

III. COMPARISON BETWEEN DATA AND SIMULATION IN ADDITIONAL RECONSTRUCTED VARIABLES

This section provides supplementary information related to Secs. VI and IX of the primary text. In this section, we provide a suite of plots comparing data and simulation for 36 different variables relevant to this analysis. These plots are shown in Figs. 5 to 35. The background fitting procedure described in Sec. VI.A.3 of the primary text is not applied here. The predictions shown in the stacked histograms in these plots come directly from the GENIE-based simulation sample. Additionally, the systematic uncertainties shown here do not include the $1\mu 1p$ CCQE constraint. In general, good agreement is observed between prediction and data across all variables.

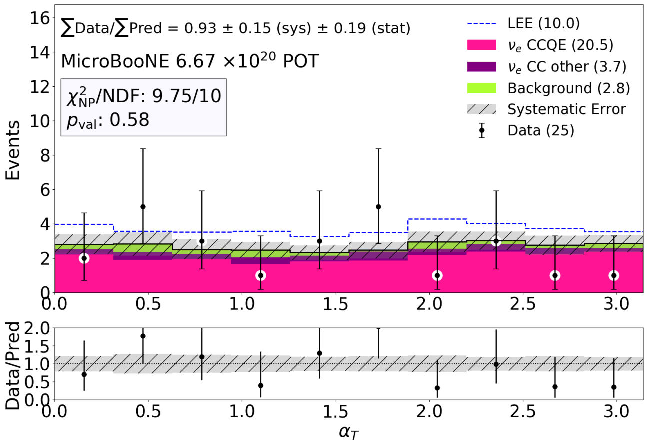

FIG. 5: Comparison between data and simulation in the $\alpha_T = \cos^{-1}\left(-\dfrac{\vec{P}^l_T \cdot \vec{P}_T}{|\vec{P}^l_T||\vec{P}_T|}\right)$ distribution.

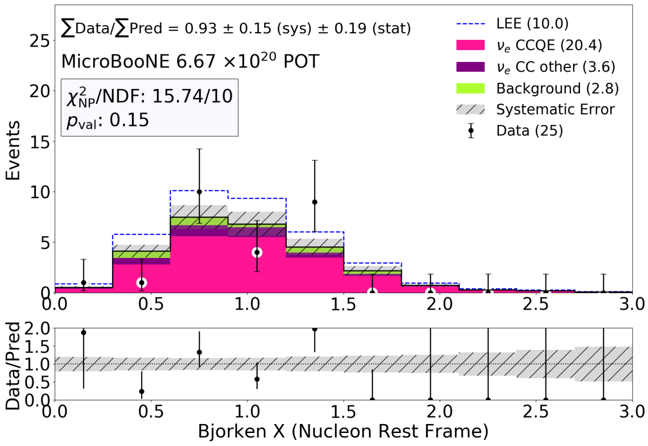

FIG. 6: Comparison between data and simulation in the Bjorken X distribution (in the nucleon rest frame).

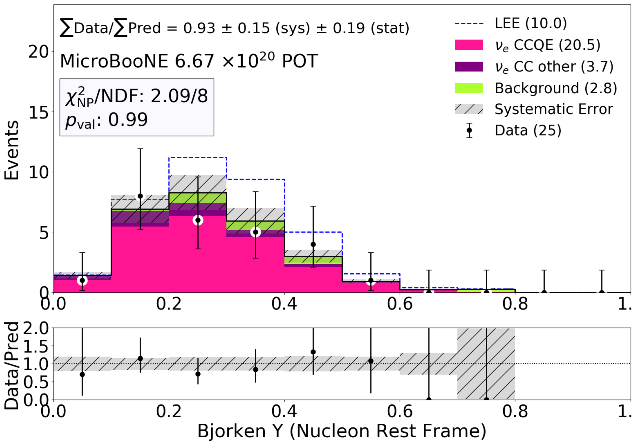

FIG. 7: Comparison between data and simulation in the Bjorken Y distribution (in the nucleon rest frame).

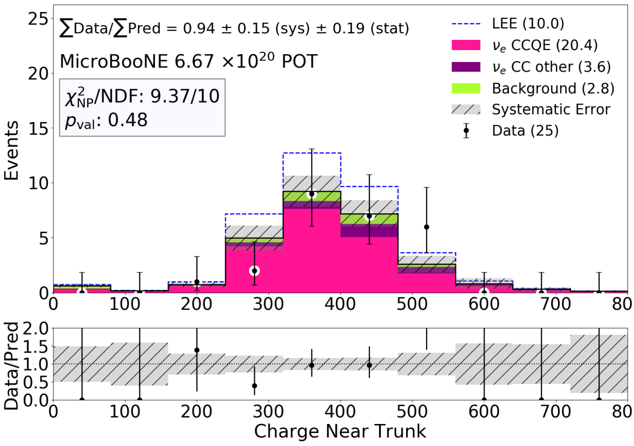

FIG. 8: Comparison between data and simulation in the distribution of charge near the trunk of the shower.

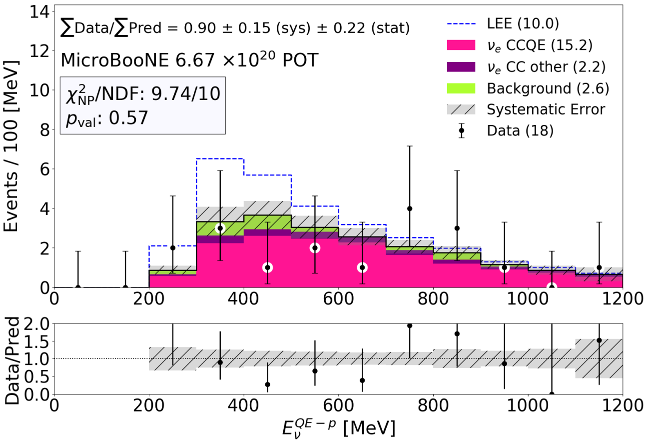

FIG. 9: Comparison between data and simulation in the E_ν^{QE-p} distribution.

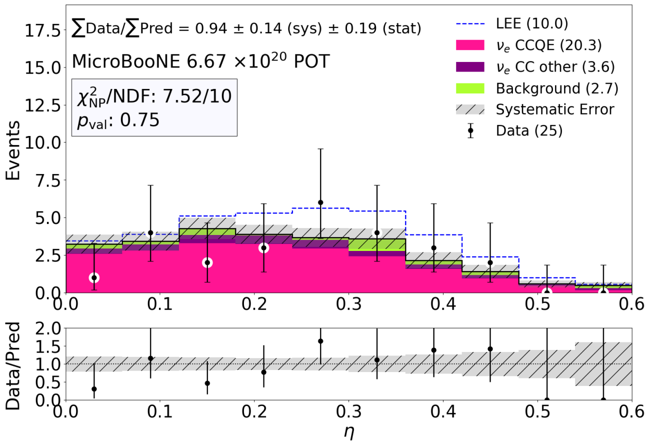

FIG. 10: Comparison between data and simulation of the ratio of difference in ionization in the trunk of particles to the sum, called the η distribution.

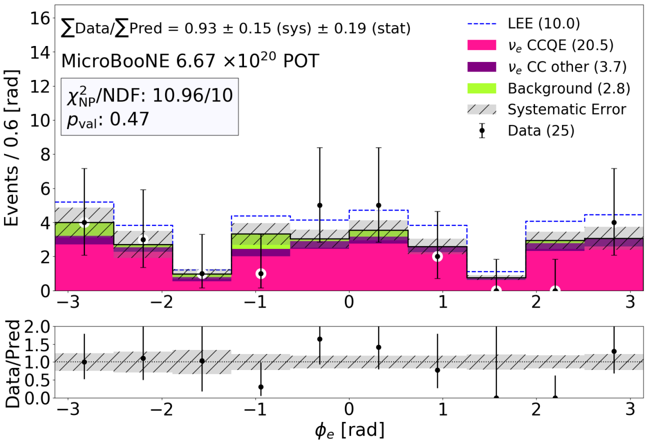

FIG. 11: Comparison between data and simulation in the ϕ_e distribution.

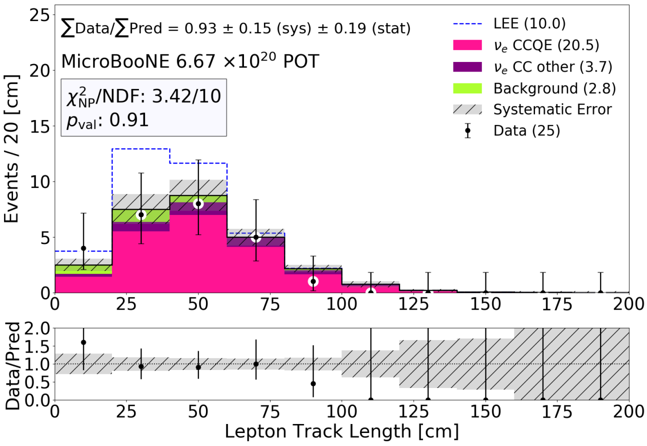

FIG. 12: Comparison between data and simulation in the lepton track length distribution.

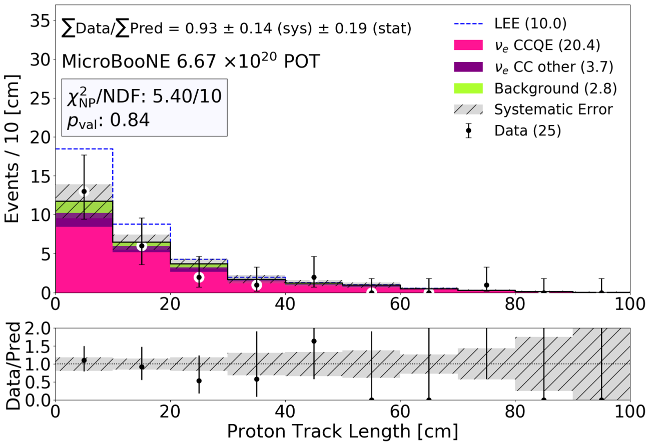

FIG. 13: Comparison between data and simulation in the proton track length distribution.

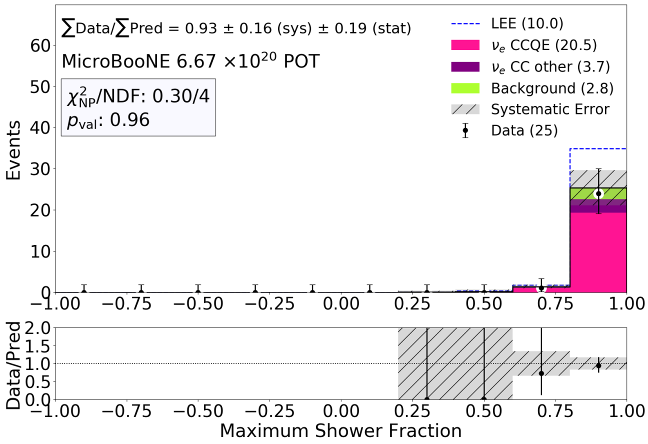

FIG. 14: Comparison between data and simulation in the maximum shower fraction distribution.

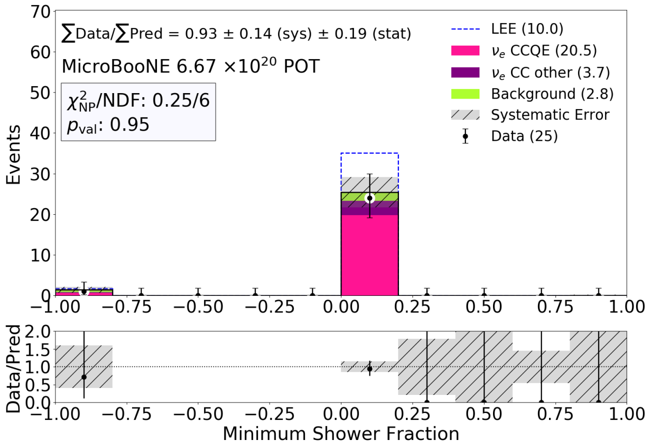

FIG. 15: Comparison between data and simulation in the minimum shower fraction distribution.

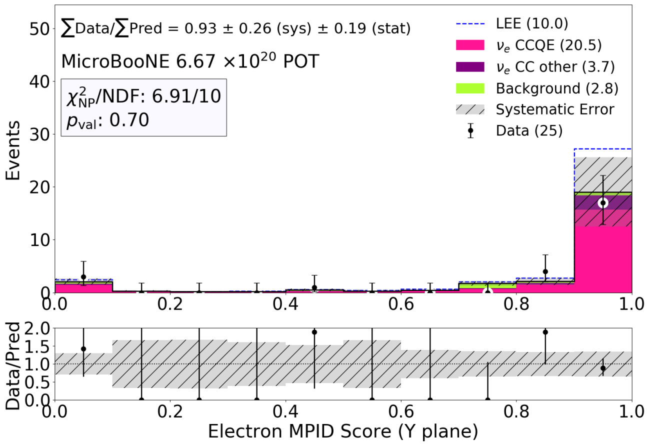

FIG. 16: Comparison between data and simulation in the distribution of the MPID electron interaction score on the Y-view plane.

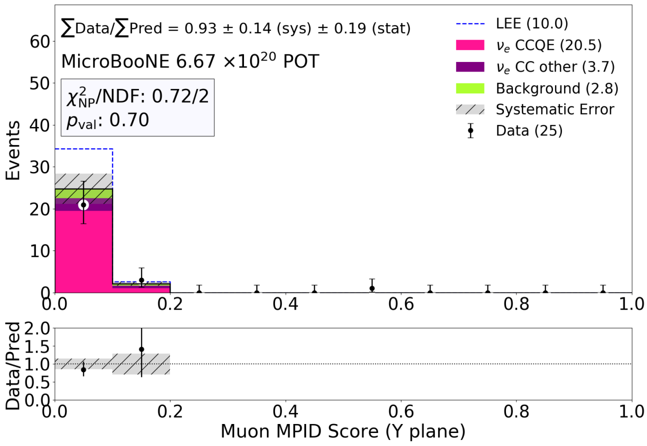

FIG. 17: Comparison between data and simulation in the distribution of the MPID muon interaction score on the Y-view plane.

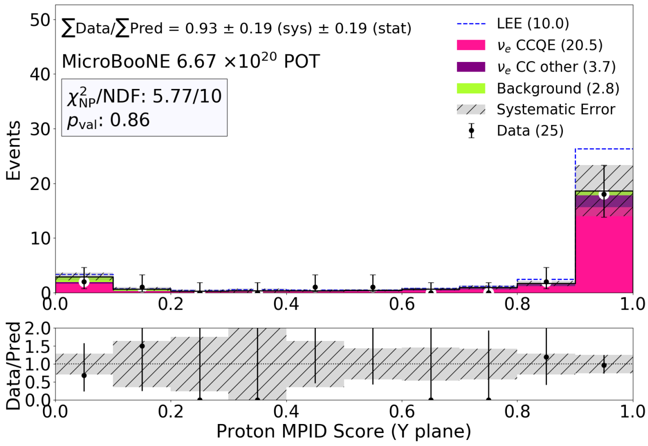

FIG. 18: Comparison between data and simulation in the distribution of the MPID proton interaction score on the Y-view plane.

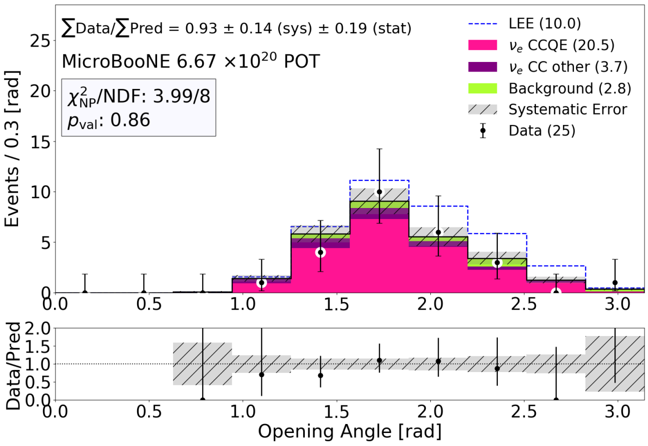

FIG. 19: Comparison between data and simulation in the opening angle distribution.

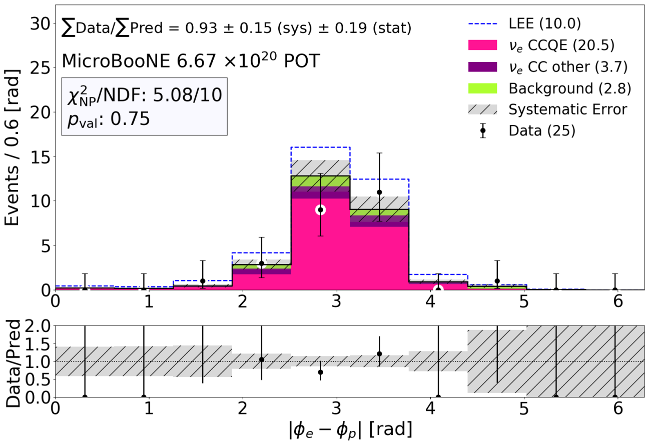

FIG. 20: Comparison between data and simulation in the $|\phi_e - \phi_p|$ distribution.

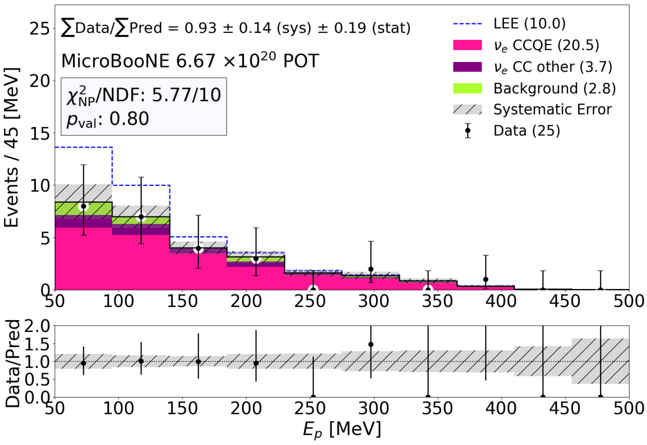

FIG. 21: Comparison between data and simulation in the E_p distribution.

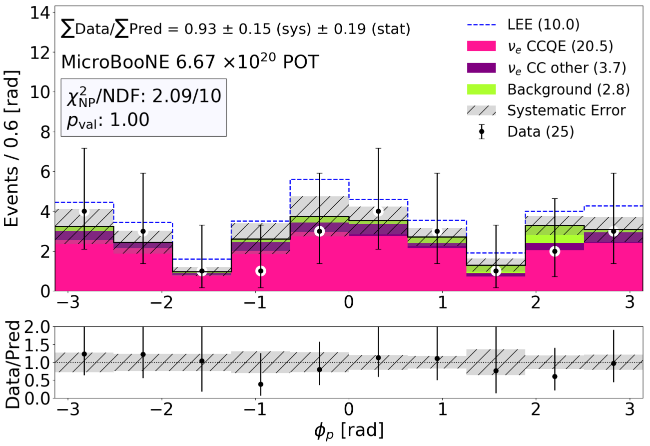

FIG. 22: Comparison between data and simulation in the ϕ_p distribution.

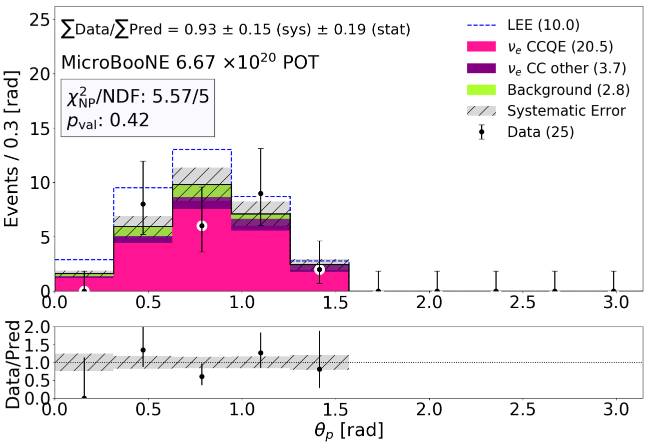

FIG. 23: Comparison between data and simulation in the θ_p distribution.

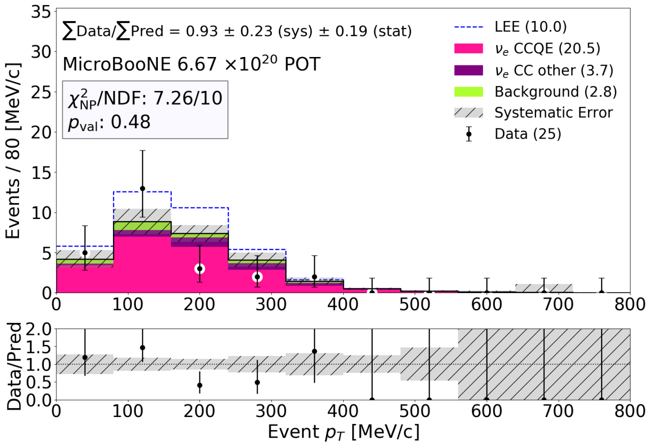

FIG. 24: Comparison between data and simulation in the total transverse momentum distribution.

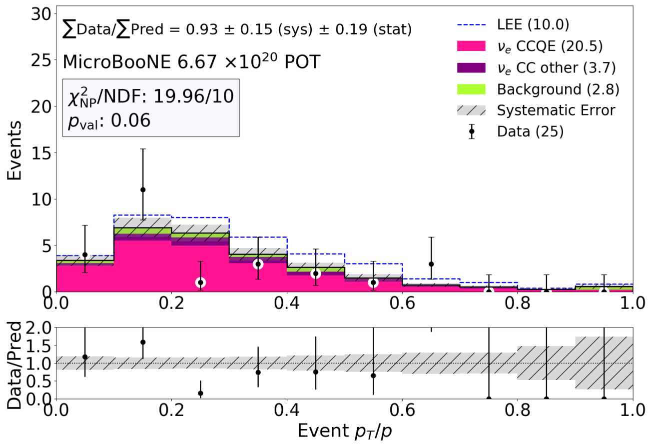

FIG. 25: Comparison between data and simulation in the transverse momentum ratio distribution.

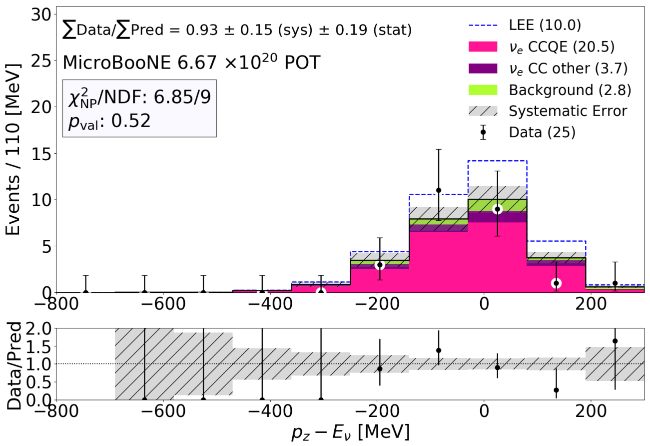

FIG. 26: Comparison between data and simulation in the $p_z - E_\nu$ distribution.

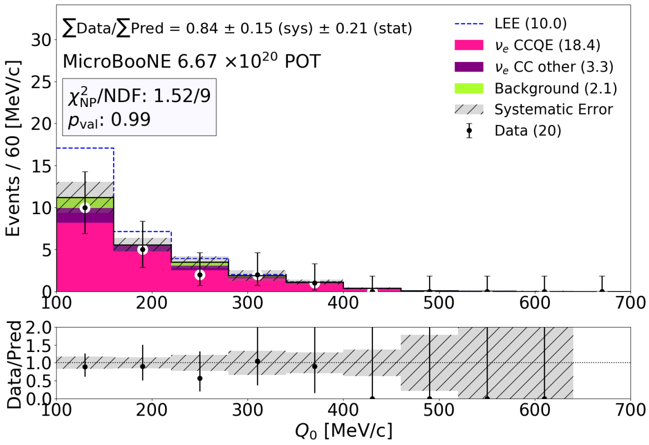

FIG. 27: Comparison between data and simulation in the Q_0 distribution.

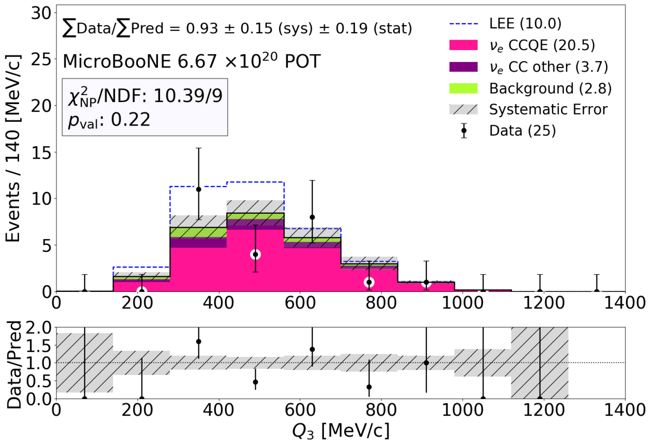

FIG. 28: Comparison between data and simulation in the Q_3 distribution.

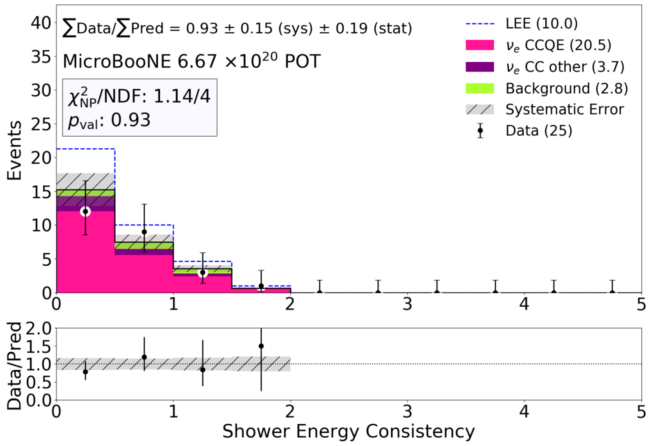

FIG. 29: Comparison between data and simulation in the shower consistency distribution.

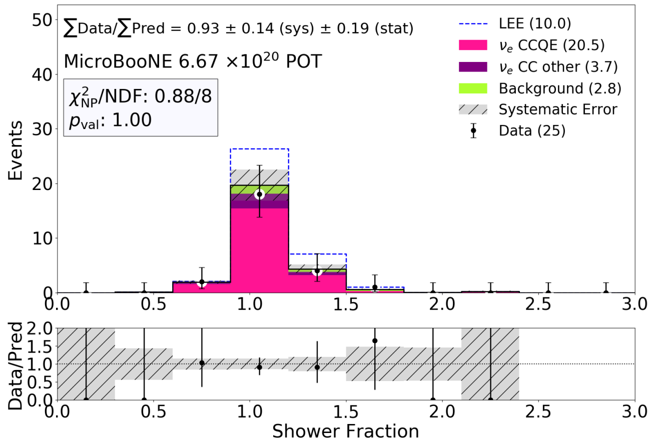

FIG. 30: Comparison between data and simulation in the shower fraction distribution.

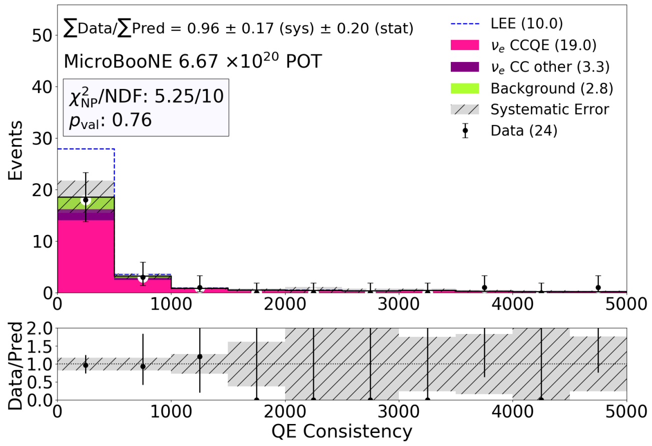

FIG. 31: Comparison between data and simulation in the quasi-elastic consistency distribution, $\sqrt{(E_\nu - E_\nu^{QE-p})^2 + (E_\nu - E_\nu^{QE-\ell})^2 + (E_\nu^{QE-\ell} - E_\nu^{QE-p})^2}$, where E_ν is energy based on range.

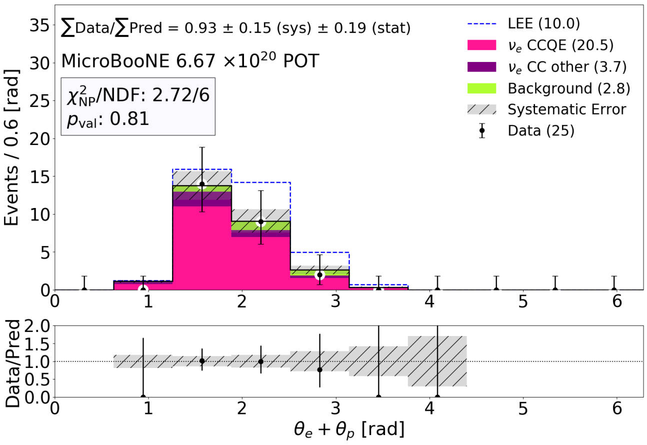

FIG. 32: Comparison between data and simulation in the $\theta_e + \theta_p$ distribution.

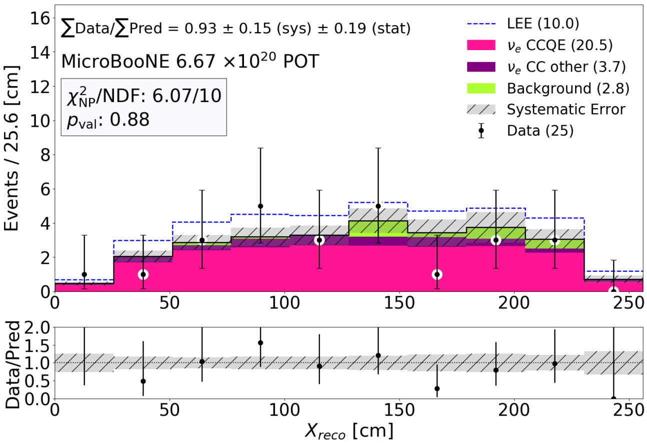

FIG. 33: Comparison between data and simulation in the X vertex position distribution.

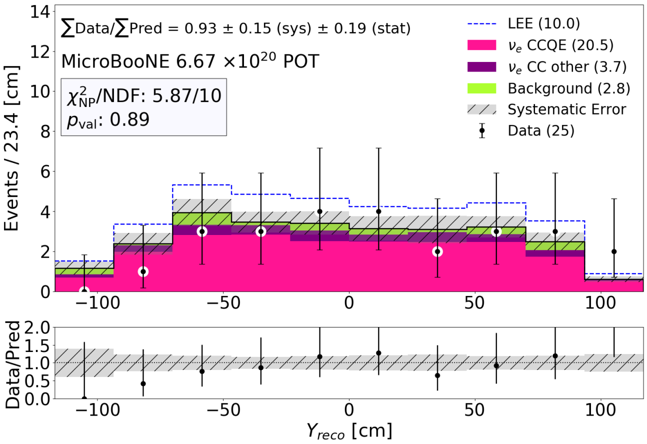

FIG. 34: Comparison between data and simulation in the Y vertex position distribution.

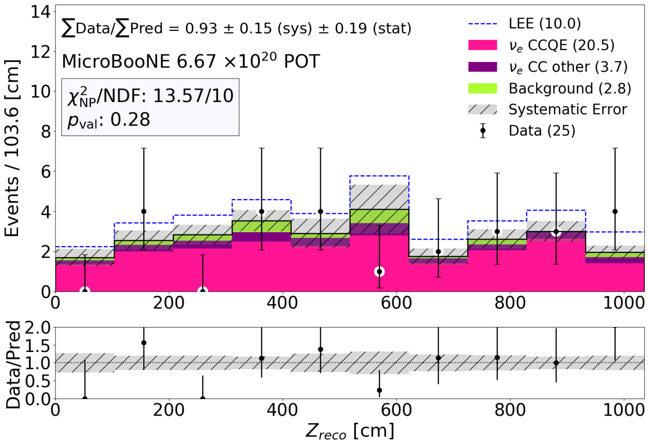

FIG. 35: Comparison between data and simulation in the Z vertex position distribution.

IV. THE $1\gamma1p$ VERTEX SAMPLE

This section provides supplementary information related to Sec. VII of the primary text. We can obtain an enhanced sample of $1\gamma1p$ vertices from the π^0 sample using a MPID muon interaction score < 0.1, thereby removing the majority of CCπ^0 events. This vertex topology matches the $1e1p$ topology and, therefore, permits a useful test of data-to-simulation agreement. First we look at the proton energy and the energy of the leading shower. The proton is the track reconstructed in the event, since events with muon tracks were largely removed. π^0 weights discussed above are applied to the MC simulation, and only MC simulation statistical errors are used. We see excellent data/prediction agreement for both the proton Fig. 36, and leading shower energy in Fig. 37.

Next, treating the leading shower as the electron, neglecting the second shower, and using the proton, the energy of the neutrino is reconstructed as though it were a $1e1p$ event. The result can be seen in Fig. 38. There is agreement across the full energy spectrum. This study gives us confidence in our ability to reconstruct the energy of low-energy $1e1p$ events.

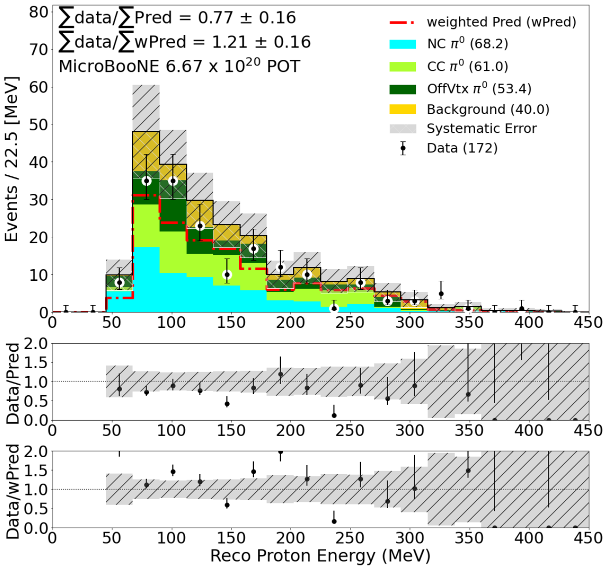

FIG. 36: The reconstructed proton energy for events passing all selection criteria (treating the event as a 1e1p event). The simulation POT has been normalized to the total POT and then re-weighted in true π^0 momentum (middle panel). In wPred, the pi0 weights have been added to the prediction (red dashed line and bottom panel).

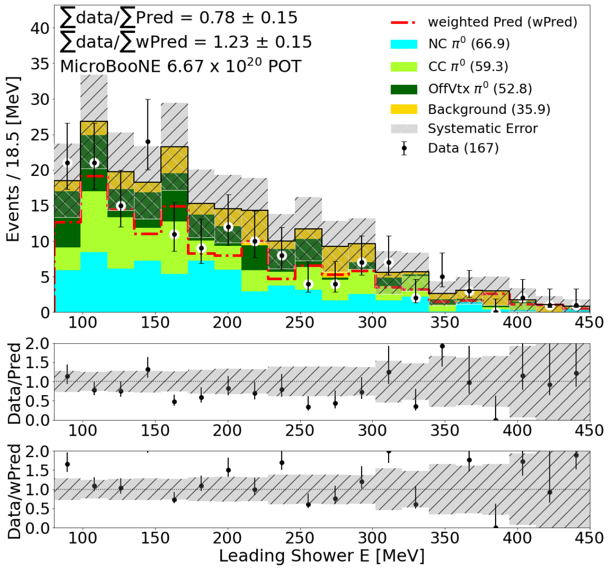

FIG. 37: The calculated leading shower energy for events passing all selection criteria (treating the event as a 1e1p event). The simulation POT has been normalized to the total POT and then re-weighted in true π^0 momentum (middle panel). In wPred, the pi0 weights have been added to the prediction (red dashed line and bottom panel).

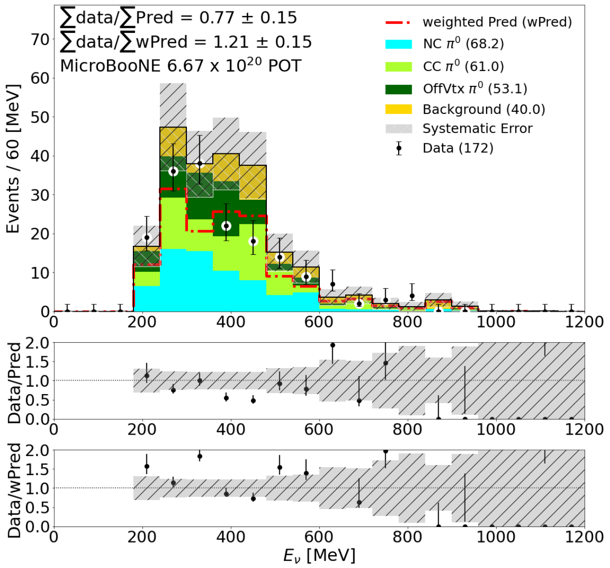

FIG. 38: The calculated neutrino energy for events passing all selection criteria (treating the event as a 1e1p event). The simulation POT has been normalized to the total POT and then re-weighted in true π^0 momentum (middle panel). In wPred, the pi0 weights have been added to the prediction (red dashed line and bottom panel).

V. ADDITIONAL COVARIANCE MATRICES

This section provides supplementary information related to Sec. VIII of the primary text. In this section, we provide a suite of plots showing additional fractional covariance and correlation matrices. In particular, Figs. 39 to 42 show the matrices coming from the flux, neutrino interaction, hadron re-interaction, and detector systematic uncertainties discussed in the text. They each show both the $1e1p$ and $1\mu1p$ analysis bins along with the covariances between them. Additionally, Fig. 43 shows the matrices for the $1e1p$ prediction after applying the $1\mu1p$ constraint.

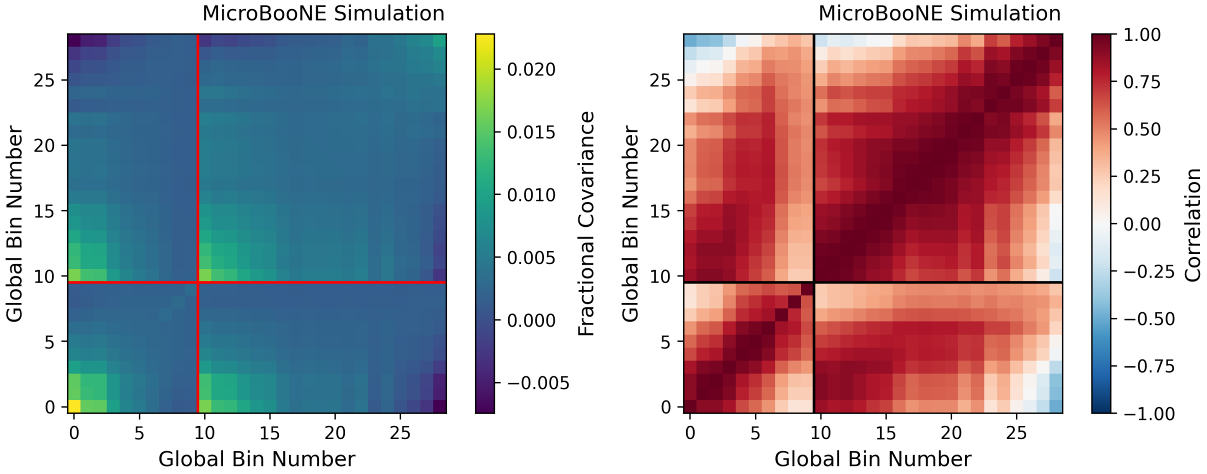

(a) Fractional covariance matrix

(b) Correlation matrix

FIG. 39: Fractional covariance (a) and correlation (b) matrices for the flux uncertainties for the $1e1p$ and $1\mu1p$ selections. The $1e1p$ events are on the lower-left and are binned from 200–1200 MeV in 100 MeV bins. The $1\mu1p$ events are on the upper-right and are binned from 250–1200 MeV in 50 MeV bins. The solid bold lines indicate the boundary between the $1e1p$ and $1\mu1p$ selections.

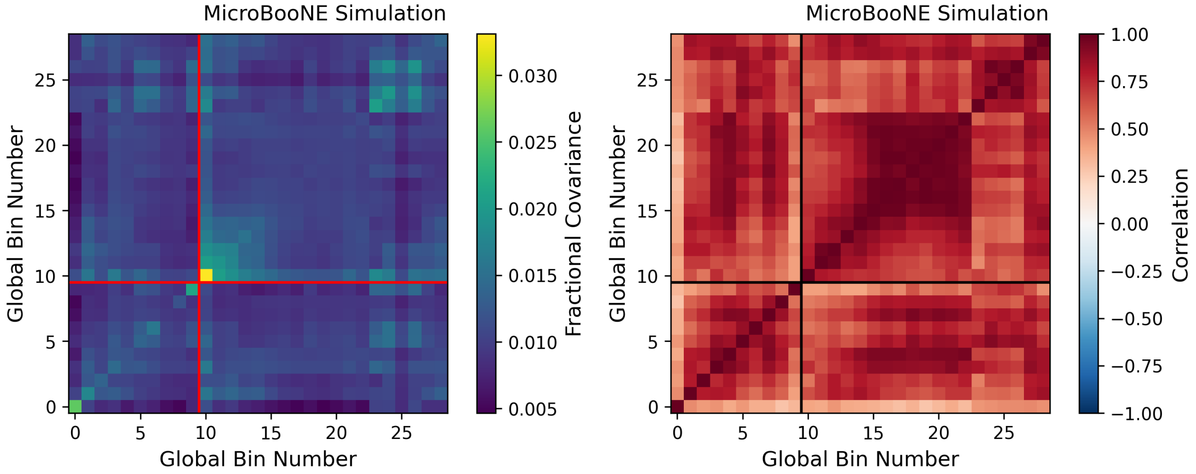

(a) Fractional covariance matrix

(b) Correlation matrix

FIG. 40: Fractional covariance (a) and correlation (b) matrices for neutrino interaction uncertainties for the $1e1p$ and $1\mu1p$ selections. The $1e1p$ events are on the lower-left and are binned from 200–1200 MeV in 100 MeV bins. The $1\mu1p$ events are on the upper-right and are binned from 250–1200 MeV in 50 MeV bins. The solid bold lines indicate the boundary between the $1e1p$ and $1\mu1p$ selections.

[1] J. E. Moyal, Journal of the Royal Statistical Society: Series B (Methodological) **11**, 150 (1949), https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.2517-6161.1949.tb00030.x.

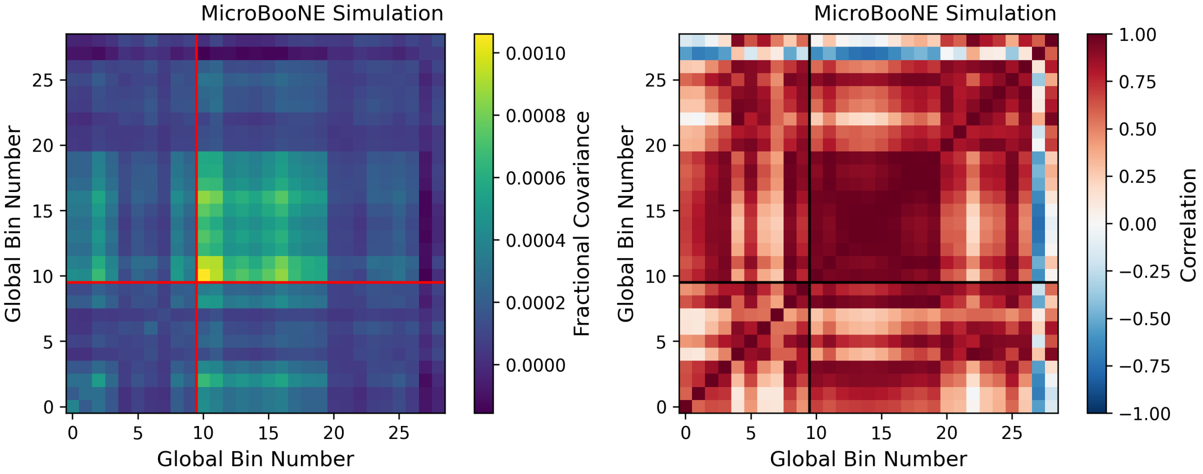

(a) Fractional covariance matrix

(b) Correlation matrix

FIG. 41: Fractional covariance (a) and correlation (b) matrices for the hadron re-interaction uncertainties for the $1e1p$ and $1\mu1p$ selections. The $1e1p$ events are on the lower-left and are binned from 200–1200 MeV in 100 MeV bins. The $1\mu1p$ events are on the upper-right and are binned from 250–1200 MeV in 50 MeV bins. The solid bold lines indicate the boundary between the $1e1p$ and $1\mu1p$ selections.

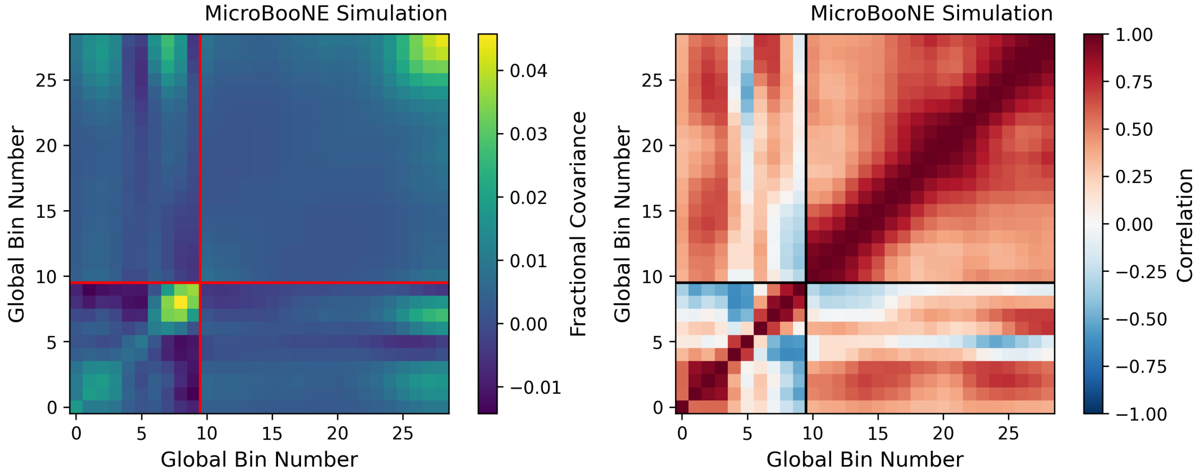

(a) Fractional covariance matrix

(b) Correlation matrix

FIG. 42: Fractional covariance (a) and correlation (b) matrices for the detector uncertainties for the $1e1p$ and $1\mu1p$ selections. The $1e1p$ events are on the lower-left and are binned from 200–1200 MeV in 100 MeV bins. The $1\mu1p$ events are on the upper-right and are binned from 250–1200 MeV in 50 MeV bins. The solid bold lines indicate the boundary between the $1e1p$ and $1\mu1p$ selections.

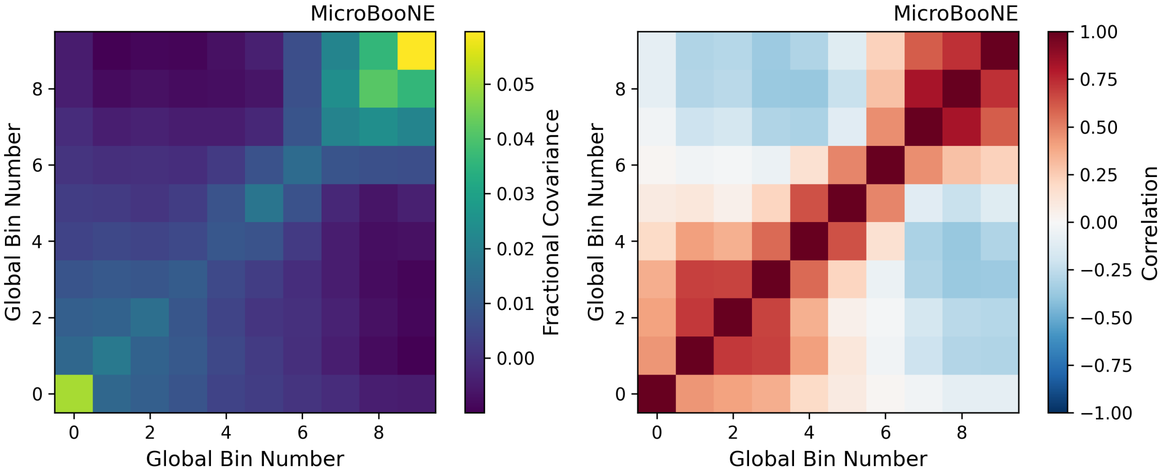

(a) Fractional covariance matrix

(b) Correlation matrix

FIG. 43: Total fractional covariance (a) and correlation (b) matrices for the $1e1p$ selection after applying the $1\mu1p$ constraint. The $1e1p$ events are binned from 200–1200 MeV in 100 MeV bins.