

Boosted decision tree reweighting of simulated neutrino interactions for $\mathcal{O}(1)$ GeV neutrino cross-section measurements

Z. Lin¹, S. Akhter², Z. Ahmad Dar^{3,2}, N. S. Alex¹, M. Betancourt⁴, S. Boyd^{5,6}, H. Budd¹, G. Caceres^{7,*}, G. A. Díaz^{4,1}, J. Felix⁸, L. Fields⁹, A. M. Gago¹⁰, P. K. Gaur², S. M. Gilligan¹¹, R. Gran¹², D. A. Harris^{13,4}, A. L. Hart¹⁴, J. Kleykamp^{1,†}, A. Klustová¹⁵, D. Last^{1,16}, A. Lozano^{7,‡}, X.-G. Lu^{5,17}, S. Manly¹, W. A. Mann¹⁸, K. S. McFarland¹, O. Moreno^{3,8}, J. K. Nelson³, V. Paolone⁶, G. N. Perdue^{4,1}, C. Pernas³, M. A. Ramírez^{16,8}, N. Roy¹³, D. Ruterbories¹, H. Schellman¹¹, C. J. Solano Salinas¹⁹, D. S. Correia⁷, M. Sultana¹, N. H. Vaughan¹¹, A. V. Waldron^{14,15}, B. Yaeggy^{20,§} and L. Zazueta^{3,||}

(MINERvA Collaboration)

¹*Department of Physics and Astronomy, University of Rochester, Rochester, New York 14627, USA*

²*Department of Physics, Aligarh Muslim University, Aligarh, Uttar Pradesh 202002, India*

³*Department of Physics, William & Mary, Williamsburg, Virginia 23187, USA*

⁴*Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

⁵*Department of Physics, University of Warwick, Coventry, CV4 7AL, United Kingdom*

⁶*Department of Physics and Astronomy, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, USA*

⁷*Centro Brasileiro de Pesquisas Físicas, Rua Dr. Xavier Sigaud 150, Urca, Rio de Janeiro, Rio de Janeiro, 22290-180, Brazil*

⁸*Campus León y Campus Guanajuato, Universidad de Guanajuato, Lascruain de Retana No. 5, Colonia Centro, Guanajuato 36000, Guanajuato México*

⁹*Department of Physics and Astronomy, University of Notre Dame, Notre Dame, Indiana 46556, USA*

¹⁰*Sección Física, Departamento de Ciencias, Pontificia Universidad Católica del Perú, Apartado 1761, Lima, Perú*

¹¹*Department of Physics, Oregon State University, Corvallis, Oregon 97331, USA*

¹²*Department of Physics, University of Minnesota—Duluth, Duluth, Minnesota 55812, USA*

¹³*York University, Department of Physics and Astronomy, Toronto, Ontario, M3J 1P3 Canada*

¹⁴*G O Jones Building, Queen Mary University of London, 327 Mile End Road, London E1 4NS, United Kingdom*

¹⁵*The Blackett Laboratory, Imperial College London, London SW7 2BW, United Kingdom*

¹⁶*Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*

¹⁷*Oxford University, Department of Physics, Oxford, OX1 3PJ, United Kingdom*

¹⁸*Physics Department, Tufts University, Medford, Massachusetts 02155, USA*

¹⁹*Facultad de Ciencias Físicas, Universidad Nacional Mayor de San Marcos, CP 15081, Lima, Perú*

²⁰*Departamento de Física, Universidad Técnica Federico Santa María, Avenida España 1680 Casilla 110-V, Valparaíso, Chile*

 (Received 8 October 2025; revised 24 May 2026; accepted 12 June 2026; published 20 July 2026)

This paper illustrates a generic method for multidimensional reweighting of $\mathcal{O}(1)$ GeV neutrino interaction Monte Carlo samples. The reweighting is based on a boosted decision tree algorithm trained on high-dimensional space in detector final-state observables. This enables one generator's events to be reweighted so that its reconstructed particle content and kinematics distributions, as well as detector efficiency, match those of a target model. The approach establishes an efficient way to reuse legacy Monte Carlo data, avoiding regeneration. As an example, we test its use in a measurement of transverse

*Present address: Department of Physics and Astronomy, University of California at Davis, Davis, California 95616, USA.

†Present address: Department of Physics and Astronomy, University of Mississippi, Oxford, Mississippi 38677, USA.

‡Present address: Department of Physics, Drexel University, Philadelphia, Pennsylvania 19104, USA.

§Present address: Department of Physics, University of Cincinnati, Cincinnati, Ohio 45221, USA.

||Present address: Syracuse University, Syracuse, New York 13244, USA.

kinematic imbalance of the μ^- and proton in charged-current quasielastic like ν_μ events from the MINERvA experiment.

DOI: [10.1103/7zty-f2qj](https://doi.org/10.1103/7zty-f2qj)

I. INTRODUCTION AND MOTIVATION

A typical analysis to estimate neutrino interaction cross sections at ~ 1 GeV neutrino energies uses a detector simulation to measure the efficiency to reconstruct events, to leverage data constraints to predict backgrounds, and to assess the effects of detector resolution. The full Monte Carlo (MC) simulations begin with a generator, such as GENIE [1–3], NEUT [4], NuWro [5], ACHILLES [6], or GiBUU [7] to predict the particle content of a given neutrino interaction, and then simulate the response of the detectors to the particles. Typically, the second step, focused on the detector response, is far more computationally expensive than the first step of simulating the neutrino interaction. Even beyond this, regenerating the full MC simulation to assess differences in assumptions among different generators, or between their various versions, is inefficient due to the parameter estimation and random sampling process [8].

A method that would allow already simulated predictions to be reweighted to those of another generator would be computationally efficient, avoid the random sampling problem, and would allow experiments who can no longer run their detector simulation due to aging software, such as the MINERvA [9] experiment, to test the results of new simulations. However, a practical difficulty is that there are many degrees of freedom in neutrino interactions kinematically and they can produce from zero to many pions and knockout nucleons in their interactions in nuclei. In the extreme case, one generator can populate a specific final-state topology that is absent from another generator. This results in events having nonoverlapping particle content and kinematics in different neutrino generators.

This problem can be managed by not forcing identical particle content, but rather by forcing observable detector quantities to be identical. For example, stable particles whose momenta are below Cerenkov threshold are effectively invisible in a water Cerenkov detector, so there would be no need for an original “source” simulation and a desired “target” simulation to match each other in predicting those particles. But it would be important for simulations to agree on the numbers and momenta of π^+ that are produced, since those can be detected with high efficiency through the decay chain $\pi^+ \rightarrow \mu^+ \nu_\mu$, $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$. Limiting the numbers of visible kinematic quantities to match detector capabilities vastly reduces the dimensionality of the problem, at the expense of making the weighting often not directly related to the models and parameters in the simulation. For example, an event with a π^+ produced by a baryon resonance decay might be (partly) simulated by

an increased weight to events where the π^+ is produced in a final-state interaction of a single nucleon knock-out event.

This paper provides an example of such a practical, lower dimensional, detector-focused reweighting and its application to a measurement of transverse kinematic imbalance of μ^- -proton events in the MINERvA experiment. The measurement focuses on ν_μ -carbon charged-current quasielastic like (CCQE-like) events, where the event final-state content has one lepton accompanied with few knock-out nucleons and no meson. The “source” sample generated by MINERvA using GENIE v2.12.6 (v2) is reweighted into a “target” sample (with the same preset) generated by GENIE v3.04.00 (v3) AR23_20i_00_000 (AR23) tune, the latter of which now is the widely used GENIE version and tuning in the neutrino community [10,11]. After reweighting, the kinematic variable distributions of source sample match that of the target sample, where the Kolmogorov–Smirnov test score is improved closer to zero in all trained variables and untrained but correlated variables.

The determination of weights even in this lower dimensional space benefits greatly from machine-learning techniques, in this case the boosted decision tree (BDT), a tree-based method for data classification or regression. In high-energy physics, a BDT is a widely used multivariate technique, either as event classifier or as reweighter [12]. In this study, we use the BDT reweighter algorithm developed by Rogozhnikov *et al* [13,14].

Elsewhere in the literature, this BDT reweighting algorithm is employed by DUNE in evaluation of oscillation parameter biases from alternative neutrino interaction models [15]. One part of the study shows that the reweighted model predicts measurable quantities, distributions of measured Q^2 separated by pion multiplicity, and another part shows how a movable off-axis detector can be used to correct for the bias. MicroBooNE’s data driven model validation study [16] examined the latter by reweighting a model to match simulated data in inclusive CC ν_μ interactions from another prediction, where the reweighting is based on muon kinematics and true energy transfer from the neutrino to the muon. Both DUNE and MicroBooNE’s studies emphasize the poor true-to-reco energy mapping despite the good agreement in trained reweight variables. Our work has a different focus, which is demonstrating a method to reproduce directly measurable quantities; neutrino energy is not directly measurable.

A different implementation of a similar reweighting strategy was used by NOvA in the results of Ref. [17]

to adjust the behavior of their final-state interaction model (FSI) to reproduce different final-state hadron multiplicities. The variables that were tuned with the BDT reweighting were limited to this specific purpose: the true E_{lepton} and the deposited energy from protons, $\pi^\pm$, and π^0 . In contrast, our study applies the BDT reweighting in high-dimensional spaces of final-state truth quantities, aiming to transform one generator's event sample so that its reconstructed distributions will match those of a target model by modifying the distributions of true variables.

II. BOOSTED DECISION TREE REWEIGHTER

The goal of the weighting process is to find the right multiplier for each MC event so that,

$$\text{multiplier}_{\text{bin}} = \frac{w_{\text{bin,target}}}{w_{\text{bin,source}}}, \quad (1)$$

where w is the weight, and “bin” represents a single region in a multidimensional histogram of events [14]. If the number of dimensions is one or two, this ratio can be calculated easily by a direct comparison between regular histograms (1D or 2D). However, it becomes impractical for higher dimensions, such as in the application here to neutrino event generation. The CCQE-like ν_μ -carbon events for example can have 1 μ^- and multiple nucleons in the final-state, whose 3-momenta form a high-dimensional set of independent variables.

In this situation, the problem becomes one of identifying which regions of the multidimensional space are to be reweighted to get the best agreement between the weighted source distribution and the target. The BDT reweighter is designed to solve reweighting in a machine learning approach analogous to gradient boosting [18]. Datasets of the same variable spaces from source and target samples are prepared for training, which is an iterative process of building decision trees in sequence. During each iteration, a new tree is created to recursively split the source and target events into different regions at terminal nodes, also known as the leaves of the tree structure [12,19], as illustrated in Fig. 1. The splitting aims to maximize the symmetrized χ^2 defined as

$$\chi^2 = \sum_{\text{leaf}} \frac{(w_{\text{leaf,source}} - w_{\text{leaf,target}})^2}{w_{\text{leaf,source}} + w_{\text{leaf,target}}}, \quad (2)$$

where $w_{\text{leaf,source/target}}$ is the sum of weights of source / target events assigned to the leaf. χ^2 measures how much a region of source distribution is different from that of the target distribution. By maximizing χ^2 , the decision tree partitions the data into regions that are more different, therefore more relevant to reweight. At each leaf of the tree, a prediction λ_{leaf} using $w_{\text{leaf,source/target}}$ is made




FIG. 1. Example of a decision tree splitting 100 source events and 100 target events into different kinematic regions based on boolean conditions on parameters p_z^μ , p_y^μ , and T_p .

$$\lambda_{\text{leaf}} = \ln \left(\frac{w_{\text{leaf,target}}}{w_{\text{leaf,source}}} \right). \quad (3)$$

The source event with weight $w_{\text{event,source}}$ assigned to this leaf will be reweighted,

$$w_{\text{event,source}} \rightarrow w'_{\text{event,source}} = w_{\text{event,source}} \times \exp(\lambda_{\text{leaf}}), \quad (4)$$

which completes the training of one tree. Since decision trees can easily overtrain, controlling the tree's maximum depth or pruning subtrees (branches) that are deemed too specific to training sample can help avoid such effects [12].

After looping through many trees, the source distributions are reweighted gradually to agree with the target distribution. By evaluating inputs of the same variable spaces through the trees' decision chains, the reweighter is able to estimate weights for statistically independent events generated by the source generator

$$w_{\text{event,source}} \rightarrow w'_{\text{event,source}} = w_{\text{event,source}} \times \exp \left(\sum_{\text{tree}} \lambda_{\text{tree}} \right), \quad (5)$$

where λ_{tree} equals the prediction of a leaf containing this event, λ_{leaf} .

III. EVENT CATEGORIZATION

In order to apply a reweighting scheme in a fixed lower dimensional space, where a few variables will be selected from the plentiful kinematic variables in neutrino MC event final-states, event categories based on final-state topologies are introduced. In the following subsections, a reaction plane coordinate system is introduced to describe the physics picture, and a practical approach to define event

topologies and choose reweight variables for MINERvA CCQE-like events is presented.

A. Reaction plane kinematics

In this study, the neutrino event kinematics are defined in a reference frame in which the struck nuclei have zero initial momentum. A reaction plane is defined by identifying the incoming neutrino direction with the $\hat{z}$ -axis and defining the shared plane in which both the incoming neutrino and outgoing lepton three momenta to lie in the $\hat{y}$ - $\hat{z}$ plane. This reaction plane is for the hadron system, so the transverse lepton momentum direction is $-\hat{y}$ and the three-momentum transfer direction is $+\hat{y}$, as shown in Fig. 2. This can be done without loss of generality because the reaction is invariant under rotation around the $\hat{z}$ -axis, the neutrino beam direction.

The reaction plane is instructive for understanding the kinematics of lepton and knock-out nucleon in CCQE-like processes. A transverse kinematic imbalance (TKI) will appear in any nucleon momentum in the $\hat{x}$ direction and in any difference between the muon p_y and the negative of the nucleon p_y . TKI effects are caused by Fermi-motion, rescattering as hadrons leave the nucleus, and missing the momentum carried by neutrons, in addition to resolution effects, such as the MINERvA studies in [20,21]. The TKI variables ($\delta\phi_T$, δp_T , and $\delta\alpha_T$) can also be built and visualized within the reaction plane in Fig. 2.




FIG. 2. Schematic illustration of the single-transverse kinematic imbalance, $\delta\phi_T$, δp_T , and $\delta\alpha_T$, defined in the plane transverse to the neutrino direction (figure derived from Fig. 2 of Ref. [20] and Fig. 2 of Ref. [21]). The neutrino direction and $\hat{z}$ -axis is out of the page while the transverse $\hat{x}$ $\hat{y}$ is in the plane of the page. If the hadron momenta $\vec{p}_N$ (possibly a single particle) is observed, the TKI variables account for the differences between its transverse component and the true transverse momentum transfer $\vec{q}_T$.

B. Particle content and detection threshold

Event categorization is based on the visible particle content of the event. Only particles which are plausibly reconstructible by the detector are individually identified in these categories, and the kinematics of those particles may then enter into the training. Particles which are below detection threshold are not individually identified, and their kinematics are only considered in aggregate, if at all. This is done to reduce the high dimension of event final-state particle information and make the machine-learning training practical.

A particle is detectable if its kinetic energy (KE) exceeds its corresponding detection thresholds. These are, in turn, defined by the detector design and event reconstruction methods. In the MINERvA detector for example, protons may be tracked with reasonable efficiency from the interaction vertex only if they have a KE above 50 MeV. The efficiency for tracking a proton from the interaction point becomes nonzero around 50 MeV. Below that threshold, protons are only observed calorimetrically and the proton kinetic energies below threshold are a proxy for that detection. Neutrons (from the neutrino interaction or from secondary interactions) with kinetic energy above 10 MeV have a chance to be detected when they elastic scatter from hydrogen nuclei in the detector, when they knockout protons from carbon nuclei, and when those carbon nuclei deexcite producing photons or nucleons.

Names of, and above-threshold particles in, the categories used in this reweighter application are listed in Table I's left and middle column. In CCQE-like interactions of neutrinos, events with multiple neutrons over threshold are rare, and the efficiency to reconstruct a neutron increases significantly with neutron kinetic energy. Therefore as a simplification, all topologies with one or more neutrons are considered together and only the most energetic (or "leading") neutron's kinematics are included in the reweighting scheme. Similarly, the numbers of events with more than two protons above reconstruction threshold is small, and so such events are lumped into a single category, again with training based on the leading protons. In this way, the kinematic information of these final-state particles can be selected from the numerous degrees of freedom in neutrino interaction. They form a lower-dimensional parameter space that is useful for detector-focused interpretation and can be more easily reproduced by the machine learning based reweighting.

Reweighting takes place within these categories, with each category having its own BDT reweighter trained independently of the others and used to estimate weights for events of that category exclusively.

C. Reweight variables

Reweight variables are the parameters chosen to represent the defining final-state features in a detector. They are fed to reweighters for training and later for weight estimation. To capture the defining features of ν_μ -carbon CCQE-like

TABLE I. The topologies and reweight variables for CCQE-like ν_μ -carbon events. Left column: topology names. Middle column: above-threshold final-state protons/neutrons based on whether their kinetic energy exceeds the detection threshold. For MINERvA detector, thresholds are 50 MeV (proton) and 10 MeV (neutron). Right column: variables used for training and weight estimation.

Topology	Above-threshold nucleons	Reweight variables
0p0n	0 proton, 0 neutron	p_y, p_z of muon; $\sum p_x, \sum p_y, \sum p_z, \sum T_p$ over all protons
0pNn	0 proton, $N \geq 1$ neutron(s)	p_y, p_z of muon; p_x, p_y, p_z of leading neutron; $\sum p_x, \sum p_y, \sum p_z, \sum T_p$ over all protons
1p0n	1 proton, 0 neutron	p_y, p_z of muon; p_x, p_y, p_z of leading proton; $\sum T_p$ over all protons
1pNn	1 proton, $N \geq 1$ neutron(s)	p_y, p_z of muon; p_x, p_y, p_z of leading proton and neutron; $\sum T_p$ over all protons
2p0n	2 protons, 0 neutron	p_y, p_z of muon; p_x, p_y, p_z of two above-threshold protons; $\sum T_p$ over all protons
2pNn	2 protons, $N \geq 1$ neutron(s)	p_y, p_z of muon; p_x, p_y, p_z of two above-threshold protons and the leading neutron; $\sum T_p$ over all protons
Others	3 or more protons	p_y, p_z of muon; p_x, p_y, p_z of leading proton; $\sum T_p$ over all protons

topologies in MINERvA detector, we choose the momenta and calorimetric energy of above-threshold final-state particles as reweight variables. Always included as reweight variables are p_y and p_z of μ^- (lepton p_x vanishes in the reaction plane coordinates, by construction) and calorimetric energy, $\sum T_p$ the sum of the kinetic energy of all protons. If there are 1 or 2 above-threshold protons, p_x, p_y , and p_z of these protons are also included; if there are 3 or more above-threshold protons, only p_x, p_y , and p_z of the leading proton are included; and if there is no above-threshold proton, the momenta $\sum p_x, \sum p_y$, and $\sum p_z$ summed over below-threshold protons are included. If there are 1 or more above-threshold neutron(s), only the leading neutron's p_x, p_y and p_z are included. These reweight variables are summarized in Table I's right column.

The above choices are optimized for the goals of this study, which are to reproduce variables of interest that are reconstructable in the MINERvA detector and used for its physics results. If the goal were to study multineutron production in neutrino interactions, for example, a different set of categories might be chosen. The method could also be adopted to event types beyond CCQE-like, for example, the charged-current events with 1 pion in the final-state ($\text{CC}1\pi$), for which the reweight variables may include the momenta of above-threshold final-state pion. A general approach would: (1) categorize the events based on the final-state particle content of interest and detection thresholds, (2) choose reweight variables to reflect features seen in the detector, and (3) train reweighters within each category to determine weights for events.

IV. REWEIGHT RESULTS

This section presents the results of reweighting MINERvA medium-energy (ME) ν_μ -carbon CCQE-like events from source generator GENIE v2.12.6 to target generator GENIE v3.04.00 AR23_20i_00_000. GENIE v2 is chosen as the source because MINERvA's large sample of simulated events was generated with this version. The neutrino events are analyzed through ROOT's Python

interface [22,23] and NUISANCE event record format [24]. Four million event source and target samples were generated using the MINERvA “medium-energy” forward-horn current (neutrino dominant) beam [25,26]. The GENIE v2 samples were preprocessed to remove a category of unphysical final-state interaction (FSI) events, see discussion at [27]. However, they were not replaced with no-FSI events as is done in the usual MINERvA analyses; instead the BDT is deciding what should be done. To ensure that source and target model share the same phase space, the 2p2h events with $q_3 > 1.2$ GeV/c in GENIE v3 are dropped, as this phase region is cut out in GENIE v2's default implementation of Valencia 2p2h model [28].¹ CCQE-like events from source and target samples were selected and divided into subsamples according to the seven categories listed, and seven reweighters are trained independently within the categories using corresponding reweight variables. The reweighter architecture's number of trees, depth of trees, and minimum sample at leaves are tuned for improving agreement with target MC for each category. For category 1p0n, ensemble of 20 trees with maximum depth of 30 are trained. For all other categories, ensembles of 100 trees with maximum depth of 4 are trained. (The need for a different architecture for the 1p0n sample is an empirical finding: the default architecture fails to reproduce the target sample as measured by the metrics in the following section.) All training requires at least 30 events at each leaf. The source and target models have different predictions for cross sections in each of the categories, as shown in Fig. 3. Reweighters modify the shape of distributions, but will not match the total cross section of each category. A normalization constant per category is calculated to change the total cross section of a

¹Upon request, MINERvA can supply a version of GENIE v2 for use with events generated and distributed with our open data product. It has all the functionality of the original GENIE v2 code, plus options to turn on bug fixes and back ported 2p2h functionality from later versions of GENIE. The code and its dependencies have been modified to build on c.2025 era software platforms.




FIG. 3. Categorical histogram of MINERvA ME CCQE-like ν_μ -carbon cross section contributed from the 7 categories listed in Table I. Orange: GENIE v2.12.6. Blue: GENIE v3.04.00 AR23_20i_00_000.

given category from the statistically independent GENIE v2 test sample prediction to that of the target. Categories are then combined to make a total prediction for observables.

A. Reweighting performance with combined categories

For evaluation of the performance, the resulting weights are applied to a statistically independent sample of GENIE v2.12.6 then compared to a sample from the target GENIE v3.04.00 AR23_20i_00_000. This subsection discusses the performance for the reweighted categories after they are combined. The individual results for each category can be found in Figs. 12–18 of Appendix A.

The combined differential cross sections of categories with 1 or more above-threshold proton(s) (“1p0n,” “1pNn,” “2p0n,” “2pNn,” “others”) in the final-state are shown in Fig. 4; reweight variables (leading proton’s momenta, proton calorimetric energy, and μ^- momenta) along with untrained observables such as TKI variables $\delta\alpha_T$, δp_T , $\delta\phi_T$, leading proton kinetic energy T_p and angle θ are plotted. The combined differential cross sections of all categories (including ones without protons, “0p0n” and “0pNn”) are shown in Fig. 5. Only the trained reweight variables are plotted, though the proton calorimetric momenta and energy are the sum of separately trained calorimetric low KE and tracked high KE components. Frequency histograms are converted to differential cross sections, where subsample cross sections of the source test sample are scaled to match that of the target sample, so the total cross sections combined are matched in magnitude.

As shown in Fig. 4, reweighters are able to identify the differences across source and target samples, which are mostly compensated by the source sample’s new weights. The two-sample Kolmogorov-Smirnov (K-S) test statistic D_{KS} is defined as the maximum distance between two empirical cumulative distribution functions $F_m(x)$ and $F_n(x)$,

$$D_{KS} = \max_x |F_m(x) - F_n(x)|, \quad (6)$$

where m and n are the samples being compared [29].² D_{KS} before and after reweighting is printed on the subplots of Fig. 4 for each kinematic variable. The expected distribution of the K-S test statistic under the null hypothesis, that the two samples follow the same underlying distribution, depends on the sample size and the distribution of weights in the sample. Accordingly, it is difficult to use the K-S test to assess compatibility with the null hypothesis after weighting, but we can use it to demonstrate the improvement in agreement between the two samples by the reduction of the K-S statistic. The reweight effects are also transferred to the variables derived from μ^- and leading proton momenta, such as TKI variables, although they are not part of the training process. This demonstrates the correlations among the target variables in the training are correctly reproduced.

Ratios of source distributions comparing to target distributions before and after reweight are plotted for reweight variables and derived quantities. In most kinematic regions, these ratios tend to 1 after reweight. This is obtained even for the TKI variables. In Fig. 5, dramatic differences are observed in calorimetric momenta and energy, where spikes at zero are present in source test sample but not in target sample. They correspond to zero KE proton events coming from “0p0n” and “0pNn” categories (see their cross sections in Appendix A Figs. 12 and 13). The reweighters remove the spikes by assigning zero weights to such events, and a large number of these zero weights can be seen in these categories in Figs. 12 and 13. These zero weights reflect the reweighter’s decision tree aggressively maximizing the symmetrized χ^2 defined in Eq. (2) to address the differences in the proton distributions.

Details about the mechanism that caused these spikes seen in GENIE v2 are reported in Appendix B. There are numerous other differences between GENIE v2 and v3, more than can be described here, and the BDT is easily able to take care of them.

B. Reweight training using smaller set of categories and variables

In this subsection, the impact of choices of categories and reweight training variables is explored, specifically the case where they contain less information about the events than the default choices listed in Table I. We look at our ability to reproduce two test measurements: muon + proton TKI variables for events with a trackable proton, and visible proton energy for all events. In a four step study, we first try a single reweighter trained on all CCQE-like events (single category) using only muon kinematics (p_y , p_z). Second,

²Reference [30] shows an example of usage in nuclear physics, where K-S test is applied to rapidity distributions to identify nuclear fragmentation processes.




FIG. 4. Differential cross sections of categories “1p0n,” “1pNn,” “2p0n,” “2pNn,” and “others” combined are plotted with respect to leading proton p_x , p_y , p_z (a, b, c); calorimetric energy $\sum T_p$ (d); $\mu^- p_y$, p_z (e, f); TKI variables δp_T , $\delta\alpha_T$, $\delta\phi_T$ (g, h, i); and leading proton T_p , θ (j, k). A frequency histogram of weights (l) is also shown. Error bars (visible only in the ratios) are statistical only. Green: test sample GENIE v2.12.6 (v2). Blue: reweighted test sample (v2'). Red: target sample GENIE v3.04.00 AR23_20i_00_000 (v3). cross section ratios of v2 and v2' comparing to v3 are plotted under each histogram, in yellow and purple respectively. K-S test statistic D_{KS} before (v2 comparing to v3) and after (v2' comparing to v3) reweighting is printed on each histogram.

with the same single category, the sum of proton kinetic energy is included in training. The sum of kinetic energy is an observable in the detector, but is not directly used in forming TKI variables. In the third and fourth scenarios, the single category is divided into two samples, one with no above-threshold protons and the other with one or more protons above-threshold. Two reweighters are trained separately using muon kinematics and sum of proton kinetic energy ($\sum_{\text{proton}} T_p$). Finally in the fourth scenario, leading proton kinematics (p_x , p_y , p_z) are also included in reweight training.

The differential cross section with respect to TKI variables are plotted in Fig. 6. As expected, no visible improvements are seen in top panel (a, b, c), since proton

kinematics are never part of the reweight. D_{KS} score, which can be interpreted as the absolute value of maximal difference between two distributions, are reduced gradually as we add subsamples and additional training variables to the reweighter from top to bottom panels of Fig. 6. When leading proton kinematics are included in training (j, k, l), all three distributions are improved significantly, especially in low $\delta\phi_T$ region where v2' almost overlaps with v3. This is a direct consequence of training on proton kinematics, because leading proton tends to travel in the opposite of muon transverse direction and events of such populate low $\delta\phi_T$ region. Similar improvements are observed for the $\sum T_p$ differential cross section, as shown in Fig. 7.

This study shows that it is important to separate out the individual trackable proton category and train on the variable




FIG. 5. Differential cross sections of all categories combined are plotted with respect to calorimetric momenta $\sum p_x$, $\sum p_y$, $\sum p_z$, and energy $\sum T_p$ summed over all final-state protons (a, b, c, d); and $\mu^- p_y$, p_z (e, f). A frequency histogram of weights (g) is also shown. Error bars (visible only in the ratios) are statistical only. Green: test sample GENIE v2.12.6 (v2). Blue: reweighted test sample (v2'). Red: target sample GENIE v3.04.00 AR23_20i_00_000 (v3). cross section ratios of v2 and v2' comparing to v3 are plotted under each histogram, in yellow and purple respectively. K-S test statistic D_{KS} before (v2 comparing to v3) and after (v2' comparing to v3) reweighting is printed on each histogram.

of interest for total visible proton energy to produce matched TKI distributions. Note that none of the more limited training scenarios achieve the quality of the results presented in the full division of categories in Fig. 4, which add more subsamples with subleading protons and neutrons, and more kinematic information about these particles. The conclusion we draw from this exploration is that it is important to validate the performance of the reweighter on the desired target distribution, particularly when it is not directly used in the training sample.

C. Application: Efficiency calculations for cross section measurements

In this subsection, the detector efficiency of reconstructed events is measured from the test sample, the reweighted test sample, and the target sample to demonstrate the reweighter's ability to recreate the target model's prediction of detection efficiency. Since the MINERvA detector uses planar targets and scintillator modules, it is better at measuring particles with high energy that travel along the beam axis than those which travel transverse to the beam

direction. This effect must be accounted for in the proton tracking efficiency model which is an input to a cross section measurement. Correcting for the efficiency is a standard step in producing a cross section. When efficiencies are small, the fractional uncertainties are large and almost always must be estimated with the help of a full simulation. The correlations between measured variables of interest and the angles of produced particles with respect to the beam depend on the model, and are different between the source and target model in this study. In this study, we use a toy model of the efficiency, ϵ , as a function of the leading proton angle with respect to the neutrino direction, θ_p , and its kinetic energy, T_p . This toy model, $\epsilon(\theta_p, T_p)$ is given by

$$\epsilon(\theta_p, T_p) = \min \left(\max \left(\frac{T_p \cos(\theta_p) - 60 \text{ MeV}}{60 \text{ MeV}}, 0.0 \right), 1 \right), \quad (7)$$

where 60 MeV was chosen above the KE selection threshold to ensure that the reweighting effectively describes the target




FIG. 6. Differential cross sections with respect to TKI variables using alternative training schemes with less complete sets of training variables and categories. (a, b, c): single CCQE category trained on muon p_y , p_z . (d, e, f): single CCQE category trained on muon p_y , p_z , and $\sum_{\text{proton}} T_p$. (g, h, i): the single CCQE category is broken into those with no above-threshold proton and those with ≥ 1 above-threshold protons, trained on muon p_y , p_z and $\sum T_p$. (j, k, l): same categories as g, h, i, but leading proton p_x , p_y , p_z are also trained in reweight. Only events with one or more above-threshold proton(s) are plotted for these TKI variables.




FIG. 7. Differential cross sections with respect to $\sum_{\text{proton}} T_p$ using alternative training schemes with less complete sets of training variables and categories. From left to right, training variables and categories are progressively more complete as specified on subplots (a, b, c, d).




FIG. 8. Contour plot of proton detecting efficiency model $\epsilon(\theta_p, T_p)$. Efficiency ranges from 0 to 1. MINERvA has high efficiency in detecting forward traveling (low θ) energetic (high T_p) protons.

model for events whose efficiency is affected by the threshold. Events with high T_p and low θ_p have efficiency close to one in this toy model. The efficiency model is plotted for reference in Fig. 8.

To study the effect of this reweighting on the extraction of cross sections, we consider the application of this efficiency model to a two-dimensional (2D) differential cross section with respect to the muon transverse momentum and the TKI variable δp_T for events with one or more above-threshold protons for the MINERvA detector. The bin-wise differential cross sections, N , and the efficient cross section, M , are defined as,

$$N_{ij} = \sum_{\text{event}_k \text{ in bin}_{ij}} C_{ij},$$

$$M_{ij} = \sum_{\text{event}_k \text{ in bin}_{ij}} C_{ij} \times \epsilon(\theta_{p,k}, T_{p,k}), \quad (8)$$

where k is the index over simulated events, i and j index the bins in δp_T and p_μ^T , respectively, and C_{ij} is the conversion factor from event rate to flux averaged differential cross section. Note that the efficiency is a function of the events k , and not directly of i and j , so the model provides the connection between those two. The efficiency for the bin, the figure of merit in this study, is defined as their ratio,

$$\phi_{ij} = M_{ij}/N_{ij}. \quad (9)$$

In a real experiment, M_{ij} would be a measurable quantity, an event rate corrected for detector smearing, related to the true correlation between the bins of the differential cross section, and the model-dependent determination of ϕ_{ij} is used to extract the cross section M_{ij} , with systematic

uncertainties to represent uncertainties in the model's calculation of the efficiency.

The “true” differential cross section with respect to δp_T and p_μ^T , N , the efficiency weighted cross section, M , and efficiency ϕ extracted from test sample, reweighted test sample, and target sample are shown in Fig. 9. As can be seen, the efficiency in both models is lower at low p_μ^T and δp_T , but the dependence on those variables is different in the two models. Figure 10 compares the ratio of the efficiencies between the source, GENIE v2, and target, GENIE v3 AR23, models before and after the reweighting. The reweighting significantly reduces the difference between the two models, with some differences remaining at the lowest p_μ^T and δp_T where the statistics for both the efficiency weighted and true samples are very low.

D. Application: Truth-to-reconstruction migration unfolding and efficiency of the signal for an example cross section

In this subsection, we test the BDT reweighting in a realistic use case. In a real measurement using data, extracting the cross section requires several steps, including estimations of neutrino flux, signal and background processes, detector efficiency and response, unfolding to correct for smearing, and so forth. For this example, we use the BDT reweighting to change the model for the signal process in a flux-weighted measurement of δp_T from MINERvA's sample of QE-like events, and compare the resulting shifts to systematic and statistical uncertainties from MINERvA's previously published analysis of low-energy data [21]. Such a change would have little impact on the flux determination or background subtraction, but could change the efficiency because of differences in final-state kinematics, including variations of the predicted cross section with neutrino energy, and because of changes in the migration between true and reconstructed quantities. Both of these effects are predicted from the signal model, and are convolved with the detector simulation.

The simulation for this analysis is taken from the FHC (neutrino beam) medium-energy sample of ν_μ -carbon CCQE-like signal MC events from the MINERvA-Analysis-Toolkit (MAT) and MINERvA's open data product [31,32]. The neutrino interaction model in this data release is GENIE v2.12.6. We have dropped unphysical elastic FSI events, as discussed earlier. Signal events are required to have a final-state containing exactly one muon, ≥ 1 protons with $T_p \geq 50$ MeV, any number of neutrons, no photons, and no mesons. The selection on reconstructed quantities follows MINERvA's CCQE-like reco selection, as discussed in Ref. [21]. Events are divided into categories introduced in Sec. III, and truth-level quantities are fed to trained BDT reweighters that reweight these events to mimic GENIE v3's distributions.

Truth and reconstructed δp_T distributions of the selected signal events, s and r , are plotted in Fig. 11(a) using the




FIG. 9. 2D differential cross sections with respect to δp_T and p_μ^T and efficiency of bins. Top panel: true cross section N , efficient cross section M , and efficiency ϕ of test sample GENIE v2.12.6 (v2). Middle panel: the same quantities for target sample GENIE v3.04.00 AR23_20i_00_000 (v3). Bottom panel: the same quantities for GENIE v2.12.6 sample reweighted (v2').

δp_T binning from Ref. [21]. Primed quantities denote the corresponding BDT-reweighted distributions. Ratios s'/s and r'/r show that the BDT reweighting leads to large upshifts in low δp_T region and down-shifts in high δp_T region. The shifts in reconstructed space are reduced compared to those in truth space.

The reconstruction selection efficiency is defined as the number of selected events divided by total number of CCQE-like events in the bin. If truth and reconstructed bins are denoted with subscripts i and j , then δp_T reco selection efficiency φ is

$$\varphi_i = \frac{s_i}{t_i}, \quad (10)$$

where t is the δp_T distribution of all CCQE-like events. Similarly, $\varphi'_i = s'_i/t'_i$. Efficiencies φ' , φ , and the ratio φ'/φ

are plotted in 11(b). The BDT reweighting effects do change the efficiency as we might anticipate. The ratio plot in bottom panel shows that these shifts are mostly covered by the neutrino interaction fractional systematic error band of MINERvA's published δp_T cross section's (GENIE systematics) [21], except the lower δp_T bins and the last bin.

The migration matrices $\mathcal{M}$ and $\mathcal{M}'$ are 2D event-counting histograms of truth versus reconstructed δp_T bins where rows (indexed with i) are normalized. They satisfy

$$r_j = \sum_i \mathcal{M}_{ij} s_i, \quad r'_j = \sum_i \mathcal{M}'_{ij} s'_i. \quad (11)$$

$\mathcal{M}_{ij}$ and $\mathcal{M}'_{ij}$ are estimates of the probability that events from truth bin i migrate to reconstructed bin j . A bias in the result due to the change in the MC simulation is expected if the reconstructed δp_T distribution is forward-folded using




FIG. 10. 2D efficiency ratio plot. (a): ϕ_{v2}/ϕ_{v3} , efficiency of source sample GENIE v2.12.6 (v2) divided by efficiency of target sample GENIE v3.04.00 AR23_20i_00_000 (v3). (b): ϕ'_{v2}/ϕ_{v3} , efficiency of reweighted source sample (v2, primed) divided by efficiency of target sample. (c): The same ratios plotted as step functions in bins of p_μ^T . (c)'s left panel: bins of $0.1 \leq p_\mu^T < 0.8$ (GeV/c); (c)'s right panel: bins of $0.8 \leq p_\mu^T < 1.5$ (GeV/c). ϕ_{v2}/ϕ_{v3} is in yellow, and ϕ'_{v2}/ϕ_{v3} is in purple.

the correct truth information but with the wrong migration matrix. To quantify this bias, we define

$$\tilde{r}'_j = \sum_i \mathcal{M}_{ij} s'_i \quad (12)$$

as the reconstructed δp_T distributions extracted from BDT-weighted truth distribution s' and unweighted migration matrix $\mathcal{M}$. $\tilde{r}'_j$ and r'_j are plotted in Fig. 11(c), where the two curves are almost indistinguishable. A closer look at the ratio $\tilde{r}'/r'$ suggests that shifts due to wrong migration model are small and almost well covered by the neutrino

interaction systematic error band, indicating that the BDT reweighting only has small effects in truth-to-reconstructed migration.

To inspect the effects of BDT reweighting on unfolding, the migration-matrix pseudoinverses $\mathcal{M}^{-1}$ and $\mathcal{M}'^{-1}$ are derived using singular value decomposition (SVD) from $\mathcal{M}$ and $\mathcal{M}'$, with a singular value cutoff at 0.07. Note that this simplistic approach introduces statistical fluctuations and is different from MINERvA's standard cross section extraction, which uses D'Agostini's iterative unfolding method [33]. Define




FIG. 11. Truth and reconstructed (reco) δp_T distributions, reco selection efficiencies, forward-folded reco δp_T from migration matrix, and unfolded truth δp_T from migration matrix's pseudoinverse. Primed (unprimed) quantities are those of GENIE v2.12.6 with BDT reweighting turned on (off). (a): selected signal events' δp_T truth distributions, s and s' , and reco distributions, r and r' . Ratios s'/s and r'/r are plotted in bottom panel. (b): reco selection efficiency ϕ' and ϕ where t and t' are total number of events. Ratio ϕ'/ϕ is plotted in bottom panel. (c): $r' = \mathcal{M}' \cdot s'$ comparing to $\tilde{r}' = \mathcal{M} \cdot s'$, the reco δp_T forward-folded from wrong migration matrix model. $\mathcal{M}$ and $\mathcal{M}'$ are the migration matrices. Ratio $\tilde{r}'/r'$ is plotted in the bottom panel. (d): s' comparing to $u' = \mathcal{M}'^{-1} \cdot r'$ and $\tilde{u}' = \mathcal{M}^{-1} \cdot r'$, the truth δp_T unfolded from correct and wrong migration model. $\mathcal{M}^{-1}$ and $\mathcal{M}'^{-1}$ are migration matrices' pseudoinverse calculated using singular value decomposition with a cutoff at 0.07. Bias $(\tilde{u}' - u')/s'$ is plotted in the bottom panel. For all of the above quantities, error bands and bars are MC statistical errors. In addition, the fractional neutrino interaction systematic error (GENIE systematics) and the MC statistical error of MINERvA's published δp_T cross section [21] are plotted in (b), (c)'s ratio subplot and (d)'s bias subplot as gray bands.

$$u'_i = \sum_j \mathcal{M}'_{ij}^{-1} r'_j, \quad \tilde{u}'_i = \sum_j \mathcal{M}_{ij}^{-1} r'_j \quad (13)$$

as the truth δp_T distributions unfolded from r' , the weighted reconstructed δp_T , with the correct and wrong migration matrix model. u' and $\tilde{u}'$ are plotted in 11(d). A few fluctuations from truth s' are seen, as unfolding from matrices' pseudoinverse is intrinsically unstable. Bias in the unfolded truth $(\tilde{u}' - u')/s'$ is plotted in the bottom panel, where we find that these fluctuations are likely statistical. They are mostly covered by the fractional statistical error band of MINERvA's

published δp_T cross section, except in the first bin. Here the statistical band is the appropriate measure because most of the differences not captured in the forward-folding study in 11(c) are statistical fluctuations in the unfolding.

As expected, most shifts in truth-to-reconstructed migration due to the BDT reweighting as observed in either the forward-folding or unfolding studies with the wrong model are covered by the gray systematic or statistical error bands, despite a few exceptions in the efficiency ratio and unfolding bias.

We stress that the results of this study should not be interpreted as corrections to previous MINERvA

measurements of these quantities. The reference models for those measurements were not the default GENIE v2.12.6 as used in this paper, but rather tuned versions which were altered to better agree with distributions in the data.

V. CONCLUSIONS

In this study, we developed a generic method of multidimensional reweighting of generator predictions of GeV energy neutrino interaction samples that helps avoid heavy computation in MC generation and enable an efficient reuse of legacy data. The samples are reweighted using a boosted decision tree method to match quantities observable in a detector. We illustrated the method by using it to reweight the MINERvA ME CCQE-like ν_μ -carbon MC sample generated by GENIE v2.12.6 to match a sample generated by GENIE v3.04.00 AR23_20i_00_000. The reweighting divides events into categories based on observable particle multiplicity to reduce the potential number of dimensions to consider. Reweights used to estimate weights for a statistically independent GENIE v2 test sample demonstrate the ability reproduce the GENIE v3 AR23 sample's predictions, even for derived variables that were not directly part of the training. We also demonstrated that such a method can be used to reproduce measurements of efficiency as would be used in a cross-section measurement in MINERvA.

For a complete application in MINERvA or other experiments, the method could be generalized to other categories of events, such as single pion production, and to predictions generated by other neutrino interaction generators, such as NuWro or NEUT. Systematic uncertainties would need to be evaluated and implemented using the predictions of the target model, rather than the source model. This is true even in the case that the source of an uncertainty is common in the two models, such as uncertainty of the axial form factor because correlations between events in the source and target models in variables not used in the tuning may not be preserved in reweighting. Our procedure could evaluate systematic uncertainties in a target model by evaluating the changes in the target model for each uncertainty, and explicitly training the reweighting for each such variation. While the procedure would consume significantly more computing resources, we deem this approach as tractable since the training and storing of BDT reweighting examples we showed in this paper is CPU-efficient and does not rely on extensive GPU clusters.

ACKNOWLEDGMENTS

This document was prepared by members of the MINERvA Collaboration using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, Office of High Energy Physics HEP User Facility. Fermilab is managed by Fermi Forward Discovery Group, LLC, acting under Contract No. 89243024CSC000002. These resources included support

for the MINERvA construction project, and support for construction also was granted by the United States National Science Foundation under Award No. PHY-0619727 and by the University of Rochester. Support for participating scientists was provided by NSF and DOE (USA); by CAPES and CNPq (Brazil); by CoNaCyT (Mexico); by ANID PIA/APOYO AFB180002, CONICYT PIA ACT1413, and Fondecyt 3170845 and 11130133 (Chile); by CONCYTEC (Consejo Nacional de Ciencia, Tecnología e Innovación Tecnológica), DGI-PUCP (Dirección de Gestión de la Investigación—Pontificia Universidad Católica del Peru), and VRI-UNI (Vice-Rectorate for Research of National University of Engineering) (Peru); NCN Opus Grant No. 2016/21/B/ST2/01092 (Poland); by Science and Technology Facilities Council (UK); by EU Horizon 2020 Marie Skłodowska-Curie Action; by a Cottrell Postdoctoral Fellowship from the Research Corporation for Scientific Advancement; by an Imperial College London President's PhD Scholarship. We thank the MINOS Collaboration for use of its near detector data. Finally, we thank the staff of Fermilab for support of the beam line, the detector, and computing infrastructure.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: REWEIGHT RESULTS OF INDIVIDUAL CATEGORIES

This appendix shows the reweight results for events from the seven individual categories listed in Table I. Differential cross sections with respect to selected final-state kinematic variables of test sample GENIE v2.12.6, reweighted test sample, and target sample GENIE v3.04.00 AR23_20i_00_000 are plotted in Fig. 12 through 18.

For all figures in this study, GENIE v2.12.6 events with predicted weights > 100 are thrown out. We investigated in these events briefly and concluded that they occur in rare, and thus low statistics, kinematic regions, and it is therefore safe to discard them. For example, some of them are DIS events in the GENIE classification without any out going pion and so fall into our CCQE selection even though such events are unusual. To further demonstrate our point, we investigated the impact of constraining range of weights. Companion plots to Fig. 5 are produced to show the impact of discarding large weight events, where cross sections of events with weight in range $[0, 10]$, $[0, 5]$, and $[0, 3]$ are plotted in Figs. 19–21, respectively. No visible differences are seen in either cross sections or ratio plots comparing to Fig. 5, although small changes in the test statistic D_{KS} are observed.




FIG. 12. Differential cross sections of category 0p0n are plotted with respect to calorimetric momenta $\sum p_x$, $\sum p_y$, $\sum p_z$, and energy $\sum T_p$ summed over all final-state protons (a, b, c, d); and $\mu^- p_y$, p_z (e, f). Histogram of weights (g) is plotted in log scale. Green: test sample GENIE v2.12.6 (v2). Blue: reweighted test sample (v2'). Red: target sample GENIE v3.04.00 AR23_20i_00_000 (v3). cross section ratios of v2 and v2' comparing to v3 are plotted under each histogram, in yellow and purple respectively. K-S test statistic D_{KS} before (v2 comparing to v3) and after (v2' comparing to v3) reweighting is printed on each histogram.




FIG. 13. Differential cross sections of category 0pNn are plotted with respect to calorimetric momenta $\sum p_x$, $\sum p_y$, $\sum p_z$, and energy $\sum T_p$ summed over all final-state protons (a, b, c, d); $\mu^- p_y$, p_z (e, f); and leading neutron p_x , p_y , p_z (g, h, i). Histogram of weights (j) is plotted in log scale. Green: test sample GENIE v2.12.6 (v2). Blue: reweighted test sample (v2'). Red: target sample GENIE v3.04.00 AR23_20i_00_000 (v3). cross section ratios of v2 and v2' comparing to v3 are plotted under each histogram, in yellow and purple respectively. K-S test statistic D_{KS} before (v2 comparing to v3) and after (v2' comparing to v3) reweighting is printed on each histogram.




FIG. 14. Differential cross sections of category 1p0n are plotted with respect to leading proton p_x , p_y , p_z (a, b, c); calorimetric energy $\sum T_p$ (d); $\mu^- p_y$, p_z (e, f); TKI variables δp_T , $\delta \alpha_T$, $\delta \phi_T$ (g, h, i); and leading proton T_p , θ (j, k). Histogram of weights (l) is plotted in log scale. Green: test sample GENIE v2.12.6 (v2). Blue: reweighted test sample (v2'). Red: target sample GENIE v3.04.00 AR23_20i_00_000 (v3). cross section ratios of v2 and v2' comparing to v3 are plotted under each histogram, in yellow and purple respectively. K-S test statistic D_{KS} before (v2 comparing to v3) and after (v2' comparing to v3) reweighting is printed on each histogram.




FIG. 15. Differential cross sections of category 1pNn are plotted with respect to leading proton p_x , p_y , p_z (a, b, c); calorimetric energy $\sum T_p$ (d); $\mu^- p_y$, p_z (e, f); leading neutron p_x , p_y , p_z (g, h, i); TKI variables δp_T , $\delta \alpha_T$, $\delta \phi_T$ (j, k, l); and leading proton T_p , θ (m, n). Histogram of weights (o) is plotted in log scale. Green: test sample GENIE v2.12.6 (v2). Blue: reweighted test sample (v2'). Red: target sample GENIE v3.04.00 AR23_20i_00_000 (v3). cross section ratios of v2 and v2' comparing to v3 are plotted under each histogram, in yellow and purple respectively. K-S test statistic D_{KS} before (v2 comparing to v3) and after (v2' comparing to v3) reweighting is printed on each histogram.




FIG. 16. Differential cross sections of category 2p0n are plotted with respect to leading proton p_x , p_y , p_z (a, b, c); second leading proton p_x , p_y , p_z (d, e, f); calorimetric energy $\sum T_p$ (g); $\mu^- p_y$, p_z (h, i); TKI variables δp_T , $\delta \alpha_T$, $\delta \phi_T$ (j, k, l); and leading proton T_p , θ (m, n). Histogram of weights (o) is plotted in log scale. Green: test sample GENIE v2.12.6 (v2). Blue: reweighted test sample (v2'). Red: target sample GENIE v3.04.00 AR23_20i_00_000 (v3). cross section ratios of v2 and v2' comparing to v3 are plotted under each histogram, in yellow and purple respectively. K-S test statistic D_{KS} before (v2 comparing to v3) and after (v2' comparing to v3) reweighting is printed on each histogram.




FIG. 17. Differential cross sections of category 2pNn are plotted with respect to leading proton p_x , p_y , p_z (a, b, c); second leading proton p_x , p_y , p_z (d, e, f); calorimetric energy $\sum T_p$ (g); $\mu^- p_y$, p_z (h, i); leading neutron p_x , p_y , p_z (j, k, l); TKI variables δp_T , $\delta \alpha_T$, $\delta \phi_T$ (m, n, o); and leading proton T_p , θ (p, q). Histogram of weights (r) is plotted in log scale. Green: test sample GENIE v2.12.6 (v2). Blue: reweighted test sample (v2'). Red: target sample GENIE v3.04.00 AR23_20i_00_000 (v3). cross section ratios of v2 and v2' comparing to v3 are plotted under each histogram, in yellow and purple respectively. K-S test statistic D_{KS} before (v2 comparing to v3) and after (v2' comparing to v3) reweighting is printed on each histogram.




FIG. 18. Differential cross sections of category “others” are plotted with respect to leading proton p_x , p_y , p_z (a, b, c); calorimetric momenta $\sum p_x$, $\sum p_y$, $\sum p_z$, and energy $\sum T_p$ summed over all final-state protons (d, e, f, g); $\mu^- p_y$, p_z (h, i); TKI variables δp_T , $\delta \alpha_T$, $\delta \phi_T$ (j, k, l); and leading proton T_p , θ (m, n). Histogram of weights (o) is plotted in log scale. Green: test sample GENIE v2.12.6 (v2). Blue: reweighted test sample (v2'). Red: target sample GENIE v3.04.00 AR23_20i_00_000 (v3). cross section ratios of v2 and v2' comparing to v3 are plotted under each histogram, in yellow and purple respectively. K-S test statistic D_{KS} before (v2 comparing to v3) and after (v2' comparing to v3) reweighting is printed on each histogram.




FIG. 19. Companion plots of Fig. 5 in which GENIE v2.12.6 events with weight greater than 10 are dropped. Differential cross sections of all categories combined are plotted with respect to calorimetric momenta $\sum p_x$, $\sum p_y$, $\sum p_z$, and energy $\sum T_p$ summed over all final-state protons (a, b, c, d), and $\mu^- p_y$, p_z (e, f). The corresponding frequency histogram of weights is shown in (g).




FIG. 20. Companion plots of Fig. 5 in which GENIE v2.12.6 events with weight greater than 5 are dropped. Differential cross sections of all categories combined are plotted with respect to calorimetric momenta $\sum p_x$, $\sum p_y$, $\sum p_z$, and energy $\sum T_p$ summed over all final-state protons (a, b, c, d), and $\mu^- p_y$, p_z (e, f). The corresponding frequency histogram of weights is shown in (g).




FIG. 21. Companion plots of Fig. 5 in which GENIE v2.12.6 events with weight greater than 3 are dropped. Differential cross sections of all categories combined are plotted with respect to calorimetric momenta $\sum p_x$, $\sum p_y$, $\sum p_z$, and energy $\sum T_p$ summed over all final-state protons (a, b, c, d), and $\mu^- p_y$, p_z (e, f). The corresponding frequency histogram of weights is shown in (g).

APPENDIX B: GENIE PROCESSES THAT LEAD TO ZERO AVAILABLE ENERGY IN CARBON

Most of the production of zero proton kinetic energy events is from a specific combination of choices in GENIE v2. GENIE v2 subtracts 25 MeV from each nucleon when there is one or two nucleons in the final-state, after FSI if any was chosen. This accounts for the energy cost to remove these nucleons from the nucleus. The most common example for carbon is when the proton undergoing a single nucleon knockout reaction, with the resulting energy shared between two nucleons (say a pn pair) in the end. When that proton has less than 25 MeV after the sharing, the subtraction means the proton is produced exactly at rest. This happens for about 4% of QE events at MINERvA energies in GENIE v2, making a spike in the distribution. This 25 MeV subtraction is not applied in GENIE v3. There is another FSI process that divides the proton energy among three or more nucleons, with a small probability to produce only neutrons, also leading to exactly zero proton momentum and kinetic energy. This process is the same for GENIE v2 and GENIE v3, there is no 25 MeV subtraction in either case. Because this second case happens for only 0.1% of QE events, no spike is visible unless the distribution is extremely finely binned.

The “shelf” in the proton KE plot is related to this 25 MeV removal energy subtraction when there is no FSI process simulated. Such protons (in neutrino QE mode) start with at least 25 MeV based on the minimum energy transfer coded in and the required Q-value of the reaction. Then in the GENIE v2 carbon case a quanta of 25 MeV is subtracted as described above, leading to near zero kinetic energy, even with no FSI energy sharing with other nucleons. The situation is similar in GENIE v3 except for the subtraction, leading to a no-FSI spectrum with that 25 MeV of kinetic energy, while the FSI process continues to produce proton KE events down to rest. For CC interactions in nonisoscalar nuclei such as Pb, the difference between proton and neutron removal energy can induce a spike at zero KE even for GENIE v3.

In the weighting procedure, the BDT simply weights the population in the spike to zero, and weights down events in the shelf, in order to describe the GENIE v3 population. However, there are events in data down practically to zero energy, so there is limited utility in using GENIE v3 (via reweighting or directly) to describe the QE hadronic energy distribution. In some future version of GENIE, this deficiency will be solved. For MINERvA’s fully simulated GENIE v2 events, one could

imagine an *ad hoc* adjustment to the events in the shelf to now weight them down as strongly even when producing the other features of the GENIE v3 model

predictions. Such a strategy would require similarly careful interpretation of the prediction as using GENIE v3 directly.

-
- [1] C. Andreopoulos *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **614**, 87 (2010).
 - [2] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, *arXiv*: 1510.05494.
 - [3] J. Tena-Vidal *et al.* (GENIE Collaboration), *Phys. Rev. D* **106**, 112001 (2022).
 - [4] Y. Hayato and L. Pickering, *Eur. Phys. J. Spec. Top.* **230**, 4469 (2021).
 - [5] T. Golan, J. Sobczyk, and J. Żmuda, *Nucl. Phys. B, Proc. Suppl.* **229–232**, 499 (2012).
 - [6] J. Isaacson, W. I. Jay, A. Lovato, P. A. N. Machado, and N. Rocco, *Phys. Rev. D* **107**, 033007 (2023).
 - [7] O. Buss, T. Gaitanos, K. Gallmeister, H. van Hees, M. Kaskulov, O. Lalakulich, A. Larionov, T. Leitner, J. Weil, and U. Mosel, *Phys. Rep.* **512**, 1 (2012).
 - [8] L. Pickering, P. Stowell, and J. Sobczyk, *J. Phys. Conf. Ser.* **888**, 012175 (2017).
 - [9] L. Aliaga *et al.* (MINERvA Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **743**, 130 (2014).
 - [10] L. Munteanu, *Neutrino-nucleus interaction systematic uncertainties and baseline model for dune analyses*, in *Talk at CEWG Meeting Aug 15 2024, Indico* (2024), <https://indico.fnal.gov/event/65637/#3-neutrino-nucleus-interaction>.
 - [11] DUNE Collaboration (to be published).
 - [12] C. Yann, *EPJ Web Conf.* **55**, 02004 (2013).
 - [13] A. Rogozhnikov, A. Ustyuzhanin, J. Eschle, R. Lane, A. Pearce, A. Thada, K. Gizdov, K. Schubert, and S. Mejlgaard, *hep_ml v0.7.2*, Zenodo, 10.5281/zenodo.7717895 (2023).
 - [14] A. Rogozhnikov, *J. Phys. Conf. Ser.* **762**, 012036 (2016).
 - [15] A. A. e. a. Abud and O. behalf of the DUNE Collaboration, *Instruments* **5**, 31 (2021).
 - [16] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **111**, 092010 (2025).
 - [17] M. A. Acero *et al.* (The NOvA Collaboration), *Phys. Rev. D* **106**, 032004 (2022).
 - [18] J. H. Friedman, *Ann. Stat.* **29**, 1189 (2001).
 - [19] G. James, D. Witten, T. Hastie, R. Tibshirani, and J. Taylor, *An Introduction to Statistical Learning: With Applications in Python*, Springer Texts in Statistics (Springer, Cham, 2023).
 - [20] X. G. Lu, L. Pickering, S. Dolan, G. Barr, D. Coplowe, Y. Uchida, D. Wark, M. O. Wascko, A. Weber, and T. Yuan, *Phys. Rev. C* **94**, 015503 (2016).
 - [21] T. Cai *et al.* (The MINERvA Collaboration), *Phys. Rev. D* **101**, 092001 (2020).
 - [22] R. Brun and F. Rademakers, *Nucl. Instrum. Methods Phys. Res., Sect. A* **389**, 81 (1997).
 - [23] R. Brun *et al.*, root-project/root: v6.18/02, Zenodo, 10.5281/zenodo.3895860 (2019).
 - [24] P. Stowell *et al.*, *J. Instrum.* **12**, P01016 (2017).
 - [25] L. Zazueta *et al.*, *Phys. Rev. D* **107**, 012001 (2023).
 - [26] L. Aliaga *et al.* (MINERvA Collaboration), *Phys. Rev. D* **94**, 092005 (2016).
 - [27] L. A. Harewood and R. Gran, *arXiv*:1906.10576.
 - [28] J. Nieves, I. Ruiz Simo, and M. Vicente Vacas, *Phys. Rev. C* **83**, 045501 (2011).
 - [29] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing (Third Edition)* (Cambridge University Press, Cambridge, England, 2007), pp. 736–737.
 - [30] I. Melo, B. Tomášik, G. Torrieri, S. Vogel, M. Bleicher, S. Koróny, and M. c. v. Gintner, *Phys. Rev. C* **80**, 024904 (2009).
 - [31] B. Messerly, R. Fine, and A. Olivier, *EPJ Web Conf.* **251**, 03046 (2021).
 - [32] MINERvA Collaboration, MINERvA Open Data Product, <https://minerva.fnal.gov/opendata/> (2026), (accessed: 2026-04-10).
 - [33] G. D’Agostini, *Nucl. Instrum. Methods Phys. Res., Sect. A* **362**, 487 (1995).