

Vertex finding in neutrino-nucleus interaction : A Model Architecture Comparison

The MINER ν A Collaboration

F. Akbar^a A. Ghosh^{b,c} S. Young^d S. Akhter^a Z. Ahmad Dar^{e,a} V. Ansari^a M. V. Ascencio^f
M. Sajjad Athar^a A. Bodek^g J. L. Bonilla^h A. Bravarⁱ H. Budd^g G. Caceres^c T. Cai^g
M.F. Carneiro^{j,c} G.A. Díaz^g J. Felix^h L. Fields^k A. Filkins^e R. Fine^g P.K.Gaur^a R. Gran^l
D.A. Harris^{m,n} D. Jenaⁿ S. Jena^o J. Kleykamp^g A. Klustová^p D. Last^q A. Lozano^c X.-G. Lu^{r,s}
E. Maher^f S. Manly^g W.A. Mann^u K.S. McFarland^g B. Messerly^v J. Miller^b O. Moreno^{e,h}
J.G. Morfinⁿ J.K. Nelson^e C. Nguyen^w A. Olivier^g V. Paolone^v G.N. Perdue^{n,g} K.-J. Plows^s
M.A. Ramírez^{q,h} D. Ruterbories^g H. Su^v V.S. Syrotenko^u A.V. Waldron^p B. Yaeggy^b
L. Zazueta^e

^aAMU Campus, Aligarh, Uttar Pradesh 202001, India

^bDepartamento de Física, Universidad Técnica Federico Santa María, Avenida España 1680 Casilla 110-V, Valparaíso, Chile

^cCentro Brasileiro de Pesquisas Físicas, Rua Dr. Xavier Sigaud 150, Urca, Rio de Janeiro, Rio de Janeiro, 22290-180, Brazil

^dOak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA

^eDepartment of Physics, William & Mary, Williamsburg, Virginia 23187, USA

^fSección Física, Departamento de Ciencias, Pontificia Universidad Católica del Perú, Apartado 1761, Lima, Perú

^gDepartment of Physics and Astronomy, University of Rochester, Rochester, New York 14627 USA

^hCampus León y Campus Guanajuato, Universidad de Guanajuato, Lascurain de Retana No. 5, Colonia Centro, Guanajuato 36000, Guanajuato México.

ⁱUniversity of Geneva, 1211 Geneva 4, Switzerland

^jDepartment of Physics, Oregon State University, Corvallis, Oregon 97331, USA

^kDepartment of Physics, University of Notre Dame, Notre Dame, Indiana 46556, USA

^lDepartment of Physics, University of Minnesota – Duluth, Duluth, Minnesota 55812, USA

^mYork University, Department of Physics and Astronomy, Toronto, Ontario, M3J 1P3 Canada

ⁿFermi National Accelerator Laboratory, Batavia, Illinois 60510, USA

^oDepartment of Physical Sciences, IISER Mohali, Knowledge City, SAS Nagar, Mohali - 140306, Punjab, India

^pThe Blackett Laboratory, Imperial College London, London SW7 2BW, United Kingdom

^qDepartment of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104

^rCoventry CV4 7AL, UK

^sOxford University, Department of Physics, Oxford, OX1 3PJ United Kingdom

^tMassachusetts College of Liberal Arts, 375 Church Street, North Adams, MA 01247

^uPhysics Department, Tufts University, Medford, Massachusetts 02155, USA

^vDepartment of Physics and Astronomy, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, USA

^wUniversity of Florida, Department of Physics, Gainesville, FL 32611

ABSTRACT: We compare different neural network architectures for Machine Learning (ML) algorithms designed to identify the neutrino interaction vertex position in the MINERvA detector. The architectures developed and optimized by hand are compared with the architectures developed in an automated way using the package "Multi-node Evolutionary Neural Networks for Deep Learning" (MENNDL), developed at Oak Ridge National Laboratory (ORNL). Furthermore, we find that while the domain expert hand-tuned network was the best performer, the differences were negligible and the auto-generated networks performed well. There is always a trade-off between human, and computer resources for network optimization and this work suggests that automated optimization, assuming resources are available, provides a compelling way to save significant expert time.

KEYWORDS: neutrino, reconstruction, convolutional neural networks, deep learning,

Contents

1	Introduction	1
2	MINERvA dectector	2
3	Analysis Challenges in the MINERvA Experiment	3
4	Network topology of the domain-expert designed ML model	4
5	Neural architecture search (NAS)	4
6	MENNDL	5
7	Details of simulation	6
8	Results	6
9	Conclusion	8

1 Introduction

Particle physics experiments are increasing their use of Machine Learning (ML) algorithms at reconstruction level to maximize their physics output. These algorithms currently tend to use one specific architecture, which means coming up with values for the number of layers of each type and the number of nodes in each of these layers. The architecture is chosen "by hand", and then the algorithms are trained on both labelled simulated and unlabelled real particle physics data in the cases of supervised learning and unsupervised learning, respectively. The process of choosing a specific architecture is lengthy and often takes months of human effort. Though deep learning has been able to bypass the process of manual feature engineering by learning representations in conjunction with statistical models in an end-to-end fashion, neural network architectures themselves are typically designed by experts in a painstaking, ad hoc manner. Neural architecture search (NAS) has been boosted as the path forward for alleviating this pain by automatically identifying architectures that are superior to hand-designed ones.

A group from ORNL developed a package named MENNDL [1, 2] which addresses the model selection problem and eases the demands on data researchers. MENNDL leverages a large number of compute nodes, which communicate over Message Passing Interface (MPI) to distribute the task of finding the optimal architecture across the nodes of a supercomputer. In our previous paper [3], we applied ML for improving the vertex reconstruction. We developed the Convolutional-Neural network (CNN) based ML model inside the collaboration for the neutrino events using both TensorFlow and CAFFE API. We showed that the ML-based reconstruction improved the

neutrino vertex reconstruction significantly compared to traditional cut-based method. Moreover, we explored the Domain Adversarial Neural Network (DANN) based ML model developed using CAFFE API, where we used both the labelled simulated events and the unlabelled real data to train the model. Following this procedure, we were able to remove the possible biases coming from large simulated training sets.

The successful application of ML brings the following questions; is the developed algorithm the optimal solution? Can the same architecture be used for similar kinds of datasets but generating from antineutrino events? Do we need a systematic to cover the architecture choice? In order to investigate these questions, we performed a thorough comparison between an algorithm that was tuned by hand over several months of researcher's time given an intimate knowledge of a detector's performance and an algorithm that was tuned using MENNDL over a few hours with a supercomputer at Oak Ridge. Exploring many topologies for a specific physics problem may help us to understand the systematic related questions mentioned above. The performances were compared using simulated antineutrino interactions in the MINERvA Detector. These interactions provide a useful arena for a model comparison because there is a detailed hit-level simulation of the detector and an important output that is needed for future physics analyses: the spatial location of the neutrino interaction point (or "vertex") in a complex detector geometry. The comparison involves testing several different models using the MENNDL package for vertex finding, where those models were optimized over different architectures. Although automated neural architecture searches are widely used in industry, to the best of our knowledge this is the first time one is being used to analyze neutrino/antineutrino physics data.

The paper is organized as follows: In Sec. 2, we describe the MINERvA detector. Next in Sec. 3, we discuss the challenges in reconstructing the vertex of the neutrino interaction. In Sec. 4, we summarize the network topology of the artisanal model developed by the MINERvA collaboration. In Sec. 5, we discuss about NAS, followed by Sec. 6 where details of MENNDL are described. In Sec. 7 we discuss the details of the simulation, and in Sec. 8 we discuss the results. Sec. 9 provides the conclusion.

2 MINERvA dectector

The MINERvA detector [4] consists of a nuclear target region where most of the interactions considered in this paper take place, followed by a fine-grained tracker which measures the products from the upstream interactions and also serves as an active target. There is electromagnetic and hadronic calorimetry surrounding the nuclear targets and the tracking region, and a muon spectrometer (the MINOS near detector [5]) is located downstream of the hadronic calorimeter region. A right-handed Cartesian coordinate system is used in MINERvA experiment. In this co-ordinate system, the origin is located at the center of inner detector. Z-axis is along the beam line parallel to the surface of the earth. X-Y plane being orthogonal to the beam line with Y-axis pointing vertically upwards and X-axis, perpendicular to both of the other axes, horizontal pointing to beam left. In this system the beam central axis is in the Y-Z plane and points slightly downward at 3.34° . The Z-axis is chosen in a way that $Z = 1200$ cm at the front face of the MINOS near detector.

The nuclear target region of the detector is complex and consists of both active and passive elements. The region is designed to have the same target material (iron or lead) located in several different regions relative to the center of the neutrino beam and relative to the muon spectrometer which is itself highly non-symmetric with respect to the neutrino beam axis. The nuclear target design allows cross-checks of the cross section measurements between regions that have different geometric acceptances. The longitudinal segmentation of the nuclear target region is shown in Fig. 1a, where the neutrino beam is incident from the left of the figure. The transverse segmentation of the different passive nuclear targets is shown in Fig. 1b.

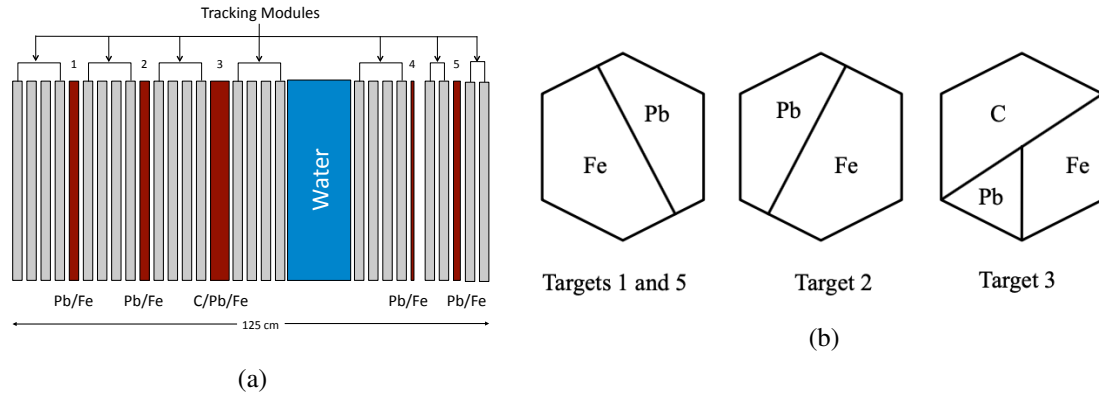


Figure 1: (a) Longitudinal segmentation (b) Transverse segmentation: targets 1, 2, 3, and 5 are shown, and target 4 consists entirely of lead. Target 3 is indicated by the figure at the right, and targets labeled 1, 2, and 5 on the left alternate between the configurations in the left and middle of the diagram at the right (figures taken from [4]).




Figure 2: Three three orientation/views of the scintillator planes.

The active scintillator planes that provide the input to the ML algorithm consist of nested bars with a triangular cross section whose base is 3.3 cm and height is 1.7 cm. The transverse cross section of the detector is hexagonal in shape, and the planes have three different orientations to allow for stereoscopic track reconstruction: scintillator bars are either vertical or they are oriented $\pm 60^\circ$ with respect to vertical. Two successive scintillator planes have strips oriented vertically followed

by one of either positive or negative 60° (Fig. 2). Each grey rectangle in Fig. 1a corresponds to two scintillator planes where the sequential non-vertical bar directions alternate between $+60^\circ$ and -60° .

3 Analysis Challenges in the MINERvA Experiment

The MINERvA experiment has a broad physics program designed to measure a wide range of neutrino interaction channels. One of the primary themes in MINERvA's physics program is to measure how different neutrino interaction channels are modified by the nuclear environment where an interaction happens. In order to do that, we use a detector that has several different thin (passive) targets made of the nuclei we want to study (for example carbon, iron, lead), and then we need to determine the target in which neutrino interaction took place. In order to determine the interaction location, segmented scintillator planes are located both upstream and downstream of the thin targets. The signals from the scintillator planes are then used to predict the start of the interaction. Neutrinos are neutral particles and leave no trace in the scintillator. Nevertheless, charged particles produced by the neutrino interactions do leave traces, emerging from the interaction vertex in all directions. In general the higher the neutrino energy, the more energy can be transferred to the nucleus, and the more energy that is transferred, the more final state charged particles can be produced. This sounds like a straightforward problem that might not require a ML algorithm, but for high-momentum transfers, there can be several charged particles leaving the interaction point in many directions. Vertex finding is made challenging by vertex occlusion in passive material, particle re-interactions, and hadronic shower cascades - these factors combine to significantly degrade vertex finding accuracy, requiring very strict sample selection cuts and background subtraction estimation procedures to accurately predict and subtract the large background. Supervised CNN could improve on the state of the art of reconstruction algorithms. In a previous publication [3] we applied Convolutional Neural Network (CNN) based supervised ML model to find the vertex. We considered specific events with high momentum transfers or in other words, events originating from deep inelastic scattering (DIS) events. Using the ML approach significantly improves the efficiency and accuracy of our vertex-finding algorithms. The CNN-based method was the top performer, including in comparisons with a tracking algorithm based on fitting clusters of hits with a Kalman filter [6]. In addition, we presented means to mitigate the possible model biases coming from the large labeled training sample using the Domain Adversarial Neural Network (DANN) [7] algorithm.

4 Network topology of the domain-expert designed ML model

Two CNN based supervised models using Caffe [8] and Tensorflow [9] application programming interfaces (APIs) were developed. Both the models are same in structure as for both API. The detailed description of the domain-expert-designed model is given in Ref. [3]. The model is indicated as Domain-expert-designed Throughout the paper. The network is made of three separate towers for X, U and V views where each tower comprises four iterations of convolution and max pooling [10][11] layers with ReLUs (Rectified Linear Unit) [12] acting as the non-linear activations. Each pooling layer consists of a kernel which decreases the dimension along the transverse axis by one. After the four iterations of convolution, ReLU and pooling, there is a fully connected layer with

196 semantic output features. The outputs for the three views are concatenated and fed to another fully connected layer with 98 outputs which in turn is input for a final fully connected layer with 67 outputs for the plane classifier, which is the input to a Softmax [13] layer. Fully connected layer allows non-linear combinations of the discovered features and operates at the end of the feature discovery layers to associate the discovered features to desired outputs. Finally, we assign the network cost using a cross-entropy function [14], which is subsequently minimized. Fig. 3 shows the structure of the Domain-expert-Designed-Tensorflow model diagrammatically. Here, the plane classifier identifies the plane where the true vertex is located. To classify the data with the plane classifier, we use a total of 67 “planecodes”. This includes two “overflow” classes - code 0 for events upstream of the detector, and code 66 for everything downstream of the last considered plane.

5 Neural architecture search (NAS)

NAS [15–18] describes a family of automated architecture optimization algorithms (i.e. optimizing the number and type of layers, hyperparameters of those layers, connectivity of layers, etc.). NAS is distinct from hyperparameter optimization which is generally restricted to optimizer algorithm choice and minor changes to layer hyperparameters of a fixed network architecture. In other words, NAS forms a scaffolding for a neural network architecture and hyperparameter search performs simple scalar adjustments or algorithm choices given an architecture. The idea of automatically learning and evolving network topologies was first explored in [19]. More recently, the pioneering works by [20] and [21] have led to a number of better, faster and cost-efficient NAS methods, thus attracting a lot of attention to this field.

There are many approaches to NAS, but they must all define a search space, search strategy, and a model evaluation strategy. The search space defines the components of the network and their possible combinations and connections. The search strategy determines the method of optimization in order to explore the search space, which significantly affects the efficiency of the search, as well as the effectiveness of the final proposed architecture. The choice of the method of optimization ensures the sufficient investigation of the search space and keeps the resulting architecture close to the optimum. Finally, the evaluation strategy compares the intermediate results, thereby helping the search strategy to choose the best option during the search process. Several techniques can be employed for the search strategies, e.g., Bayesian optimization [22], evolutionary algorithms [23], gradient descent [24], reinforcement learning [20], and meta-learning [25].

6 MENNDL

Multi-node Evolutionary Neural Networks for Deep Learning (MENNDL) [1, 2] is a NAS algorithm that utilizes evolutionary optimization for evolving the architecture and hyperparameters of convolutional neural networks on high performance computers.

MENNDL uses a master-worker paradigm such that a master process is used to perform the core evolutionary process and workers are used to evaluate the fitness of the generated networks. The evolutionary process is asynchronous. Once a sufficient number of individual networks has been evaluated, a new set of individuals are produced through selection, crossover, and mutation and then queued to be evaluated. This ensures the computational resources used by the workers are

not left idle. MENNDL has previously shown success in automatically designing neural networks for several scientific datasets [1, 2]. It is particularly useful for datasets that have very different image characteristics than the photographic imagery typically used in machine learning literature (e.g. ImageNet [26]). A key difference between MENNDL and other NAS methods is that it focuses on exploring all possible hyperparameters for the potential layers along with the composition of those layers, as opposed to most other methods which limit the range of hyperparameters to a small set (e.g. DARTS only allows for a handful of potential kernel sizes [27]).

In this work, some modifications were made in order to handle the unique data used. In order to handle the three views, the same three tower structure was used for the MENNDL networks, with each tower having identical layer types and hyperparameters learned through the evolutionary process. The maximum size of hyperparameters related to the width of the input image was set to the maximum allowed by the two smaller views. MENNDL generated and evaluated approximately 10,000 networks on the Titan Supercomputer using 500 nodes for 12 hours. Example networks are shown in Fig. 4. From these networks, a subset was chosen for further training and evaluation based on their validation accuracy.

The events used for training and developing the MENNDL networks, are the simulated events generated using neutrino flux peaked around 6 GeV.

It is noted that these models are primarily optimized over short run periods (only 5000 iterations) using the neural network software library Caffe. The version of MENNDL used in this work is defined in the 2017 paper [2], and thus the hyperparameters optimized by MENNDL in this work are limited to those defining the layers and the architecture of the neural network. We can optimize the number of layers and the type of each layer along with layer specific hyperparameters (e.g. kernel size, number of filters, dropout rate, etc.). Solver hyperparameters such as learning rate are not optimized by MENNDL.

7 Details of simulation

Among 10,000 MENNDL models, we chose nine models by looking at their validation accuracy. Validation loss and accuracy are the loss and accuracy on validation set which gives a measure of the quality of the model. Next, we considered a data sample of size 1.66 million events, and divide that sample in two: $\sim 93\%$ of total sample events are used for training whereas rest ($\sim 7\%$) are used for validation of the models. These events are the antineutrino events coming from the NuMI beam with a peak energy of around 6 GeV. We trained the MENNDL models and the domain-expert-designed model using those simulated antineutrino events up to 20 iterations (epochs). Note that, since we use a CNN-based model, we make images of the input events. We presented the data as two channel images using deposited energy and hit time to train the network. Each event contains three images - one for each view (X, U, V)- that are fed into the separate convolutional towers. We preprocessed the images by normalizing the values of the total deposited energy to unity on an event-by-event basis. Timing values were pre-processed by subtracting the event time extracted by fitting a linear model to the anchor track and scaled by the largest absolute value of all the relative times in the event. A detailed description of image processing is given in Ref. [3]. For studying the impact of the models on a physics analysis, we use another set of simulated data sample of size two million events.

8 Results

MENNDL Models	Total number of parameters	Time of execution (gpu-hour)	Validation loss after 20 epochs	Validation accuracy after 20 epochs (%)
MENNDL-1	415104	8*15.5	0.82	71.7
MENNDL-2	949248	8*23	0.87	70.1
MENNDL-3	2857568	8*16.75	1.00	67.9
MENNDL-4	2857920	8*20.5	0.85	68.5
MENNDL-5	1637712	8*20.75	0.73	74.4
MENNDL-6	1637712	8*20.75	0.74	74.2
MENNDL-7	2014224	8*20.5	0.85	70.0
MENNDL-8	3549240	8*20.5	0.94	66.2
MENNDL-9	2691616	8*20.5	0.86	70.6

Table 1: Total number of parameters and time of execution for all the caffe models in Fig.5. Time of of execution has been reported in terms of gpu-hour. We used 8 gpus (Tesla P100) and we used the same setup for all the models

Models	Validation loss after 20 epochs	Validation accuracy after 20 epochs (%)	Total number of parameters
DED-caffe	0.78	72.6	12115716
MENNDL-5	0.73	74.5	1637712

Table 2: Validation loss and accuracy numbers for each models along with the total number of adjustable parameters in each model.

The validation accuracy and loss curves of the nine chosen MENNDL models are shown in Fig. 5. Among them, we selected the model with highest validation accuracy and compared that to the domain-expert-designed-Caffe. Table 1 shows the number of trainable parameters of each of the MENNDL models and corresponding time of execution.

Fig. 6 shows the comparison of the validation accuracy and validation loss between domain-expert-designed-Caffe, and MENNDL (MENNDL-5 as indicated by the yellow line in Fig. 5) models. The red dashed line corresponds to the MENNDL model whereas the black solid line corresponds to the domain-expert-designed-Caffe model. Figs. 6a and 9b show the loss and accuracy as the number of epochs increases. The MENNDL model took ~ 20 hours to train whereas the time taken by other two models for the training is more than 24 hours. MENNDL model was automated while the others required human intervention. We can see from the plot that the MENNDL model and domain-expert-designed-Caffe model start saturating at the same number of epochs. After 10 epochs of training, both models show similar accuracy and loss. Table 2 lists the validation loss and accuracy achieved by the two models after reaching the saturation with the total number of adjustable parameters in each model. We also compared the number of parameters of each model, which is shown in the third column of Table 2. We see that models found by MENNDL have almost 10 times fewer parameters than the models designed by hand, thus having less expressivity than the more complex domain-expert-designed model. This is expected due to the process used by MENNDL, where the fitness of the model is evaluated after a limited number of iterations. However, domain-expert-designed model and MENNDL models demonstrate similar performance, validating the approach used by MENNDL [2].

Next, we test our models where we generate the ML-based predictions of the vertex of each event

using those two models followed by the incorporation of those predictions into the MINERvA analysis framework. We used a simulation sample of about 10% of the total antineutrino statistics for this purpose.

Fig. 7 shows the event distribution at target 2 with respect to vertex Z using antineutrino simulated data set in deep inelastic scattering (DIS) region. The variable vertex Z is closely related to the actual position of the nuclear targets in the MINERvA detector. We use the ML-based prediction for the reconstructed vertex position, where Figs. 7a, and 7b represent domain-expert-designed-Caffe and MENNDL models, respectively. As target 2 is composed of iron and lead the plots are made for the number of events in iron and lead separately. Here we show the event distribution generated in iron. In addition to the event distributions coming from the nuclear target (iron), the number of events coming from the plastic scintillator upstream (in this case six planes between targets 1 and 2, closest to target 2) and downstream (in this case six planes between targets 2 and 3, closest to target 2) of the nuclear target are also presented. The two event distributions are quite similar to each other which signifies that all two models provide similar predictions for the interaction vertex. In Fig.8, have took the ratio of 7b to 7a. We can see the ratio lies near 1 within error bars.

We calculate the efficiency and the purity of the sample, where efficiency is the fraction of true events reconstructed, and purity is the fraction of events where the reconstructed value was equal to the underlying true generated value. For the calculation of the efficiency, the numerator is the number of events in the source nuclear target passing the reconstructed sample selection cuts, the reconstructed and truth DIS cut. The denominator is all true generated CC DIS events in the true fiducial volume with true $\theta_\mu < 17^\circ$. Purity is related to the efficiency. The purity numerator is same as the efficiency, divided by the number of events passing the reconstructed sample selection cuts and reconstructed DIS cut.

Fig. 9 shows the efficiency and purity versus true target number. We compared two models, and the difference between them vary between 0.8% to 3.4% and from 1.0% to 2.1% for efficiency and purity, respectively (the difference is calculated between domain-expert-dedigned-Caffe and MENNDL).

9 Conclusion

In this paper, we explored the performance of a ML model obtained from the MENNDL package developed by the ORNL group using MINERvA simulated neutrino data. We further compared the MENNDL model with two successful models which were manually developed by the MINERvA collaboration for improving the vertex reconstruction precision for neutrino interactions. In this work, we trained those models using antineutrino DIS events, generated predictions for the vertex location of the antineutrino DIS events, and compared the performance of the two models in terms of vertex position accuracy at the nuclear target region, efficiency, and purity. In terms of the efficiency with respect to the true target location of DIS events, the differences between the two models ranged between 0.8% to 3.4%, and for the purity those differences varied between 1.0% to 2.1%. We show that the MENNDL generated model performs similarly to the domain-expert-designed model. We also see that the models found by MENNDL have almost 10 times fewer parameters than the domain-expert-designed model. Hence, MENNDL models have lower capacity than the comparatively more complex domain-expert-designed model. However, both models can

achieve similar performance. From this we observe that MENNDL favors lower capacity models over more complex ones. The domain-expert-designed model was carefully developed, but done manually, and therefore required a significant amount of time of persons with intimate knowledge of the detector's performance. On the other hand, the MENNDL models were generated and evaluated using the package MENNDL where the models are trained over fewer iterations, therefore significantly reducing the amount of specialized human input and time required. This technique is quite common in industry for use in automated neural architectures but to the best of our knowledge this is the first time it has been used to analyze neutrino/antineutrino physics data. In the coming years, the current running and future neutrino experiments will collect an unprecedented amount of information dense data. As a result, the implementation of a ML algorithm will become necessary to extract the physics embedded in these data, and the MENNDL package will be of immense help in significantly reducing the required human time and associated time needed to produce an optimized ML model.

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, HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359. 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 Proyecto Basal FB 0821, 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. The research here was also sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U. S. Department of Energy. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

References

- [1] S. R. Young, D. C. Rose, T. P. Karnowski, S.-H. Lim and R. M. Patton, *Optimizing deep learning hyper-parameters through an evolutionary algorithm*, in *Proceedings of the Workshop on Machine Learning in High-Performance Computing Environments*, pp. 1–5, 2015.

- [2] S. R. Young, D. C. Rose, T. Johnston, W. T. Heller, T. P. Karnowski, T. E. Potok et al., *Evolving deep networks using hpc*, in *Proceedings of the Machine Learning on HPC Environments*, pp. 1–7, (2017).
- [3] MINERvA collaboration, *Reducing model bias in a deep learning classifier using domain adversarial neural networks in the MINERvA experiment*, *JINST* **13** (2018) P11020 [[1808.08332](#)].
- [4] MINERvA collaboration, *Design, Calibration, and Performance of the MINERvA Detector*, *Nucl. Instrum. Meth. A* **743** (2014) 130 [[1305.5199](#)].
- [5] MINOS COLLABORATION collaboration, *The Magnetized steel and scintillator calorimeters of the MINOS experiment*, *Nucl. Instrum. Meth. A* **596** (2008) 190 [[0805.3170](#)].
- [6] G. Welch and G. Bishop, *An introduction to the kalman filter*, tech. rep., University of North Carolina at Chapel Hill, USA, 1995.
- [7] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette et al., *Domain-adversarial training of neural networks*, *Journal of Machine Learning Research* **17** (2016) 1.
- [8] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick et al., *Caffe: Convolutional architecture for fast feature embedding*, *arXiv preprint arXiv:1408.5093* (2014) .
- [9] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, 2015, arXiv:1603.04467.
- [10] Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard et al., *Handwritten digit recognition with a back-propagation network*, in *NIPS*, 1989.
- [11] Y. Lecun, L. Bottou, Y. Bengio and P. Haffner, *Gradient-based learning applied to document recognition*, *Proceedings of the IEEE* **86** (1998) 2278.
- [12] V. Nair and G. E. Hinton, *Rectified linear units improve restricted boltzmann machines*, in *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML'10*, (Madison, WI, USA), p. 807–814, Omnipress, 2010.
- [13] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*. MIT Press, 2016.
- [14] K. P. Murphy, *Probabilistic Machine Learning: An introduction*. MIT Press, 2022.
- [15] T. Elsken, J. H. Metzen, F. Hutter et al., *Neural architecture search: A survey.*, *J. Mach. Learn. Res.* **20** (2019) 1.
- [16] L. Li and A. Talwalkar, *Random search and reproducibility for neural architecture search*, *CoRR, Conference on Uncertainty in Artificial Intelligence(UAI), 2019* **abs/1902.07638** (2019) [[1902.07638](#)].
- [17] M. F. Kasim, D. Watson-Parris, L. Deaconu, S. Oliver, P. Hatfield, D. H. Froula et al., *Building high accuracy emulators for scientific simulations with deep neural architecture search*, *Machine Learning: Science and Technology* **3** (2021) 015013.
- [18] P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, X. Chen et al., *A comprehensive survey of neural architecture search: Challenges and solutions*, *ACM Computing Surveys (CSUR)* **54** (2021) 1.
- [19] K. O. Stanley and R. Miikkulainen, *Evolving neural networks through augmenting topologies*, *Evolutionary Computation* **10** (2002) 99.
- [20] B. Zoph and Q. V. Le, *Neural architecture search with reinforcement learning*, *CoRR* **abs/1611.01578** (2016) [[1611.01578](#)].
- [21] B. Baker, O. Gupta, N. Naik and R. Raskar, *Designing neural network architectures using reinforcement learning*, *CoRR* **abs/1611.02167** (2016) [[1611.02167](#)].

- [22] B. Ru, X. Wan, X. Dong and M. Osborne, *Interpretable neural architecture search via bayesian optimisation with weisfeiler-lehman kernels*, in *International Conference on Learning Representations*, 2021, <https://openreview.net/forum?id=j9Rv7qdXjd>.
- [23] H. Liu, K. Simonyan, O. Vinyals, C. Fernando and K. Kavukcuoglu, *Hierarchical representations for efficient architecture search*, *Accepted as a conference paper at ICLR 2018* **abs/1711.00436** (2017) [[1711.00436](https://arxiv.org/abs/1711.00436)].
- [24] S. Santra, J.-W. Hsieh and C.-F. Lin, *Gradient descent effects on differential neural architecture search: A survey*, *IEEE Access* **9** (2021) 89602.
- [25] D. Lian, Y. Zheng, Y. Xu, Y. Lu, L. Lin, P. Zhao et al., *Towards fast adaptation of neural architectures with meta learning*, in *International Conference on Learning Representations*, 2019, Published at ICLR 2019.
- [26] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li and L. Fei-Fei, *Imagenet: A large-scale hierarchical image database*, in *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255, IEEE, 2009.
- [27] H. Liu, K. Simonyan and Y. Yang, *Darts: Differentiable architecture search*, *arXiv preprint arXiv:1806.09055* (2018) .

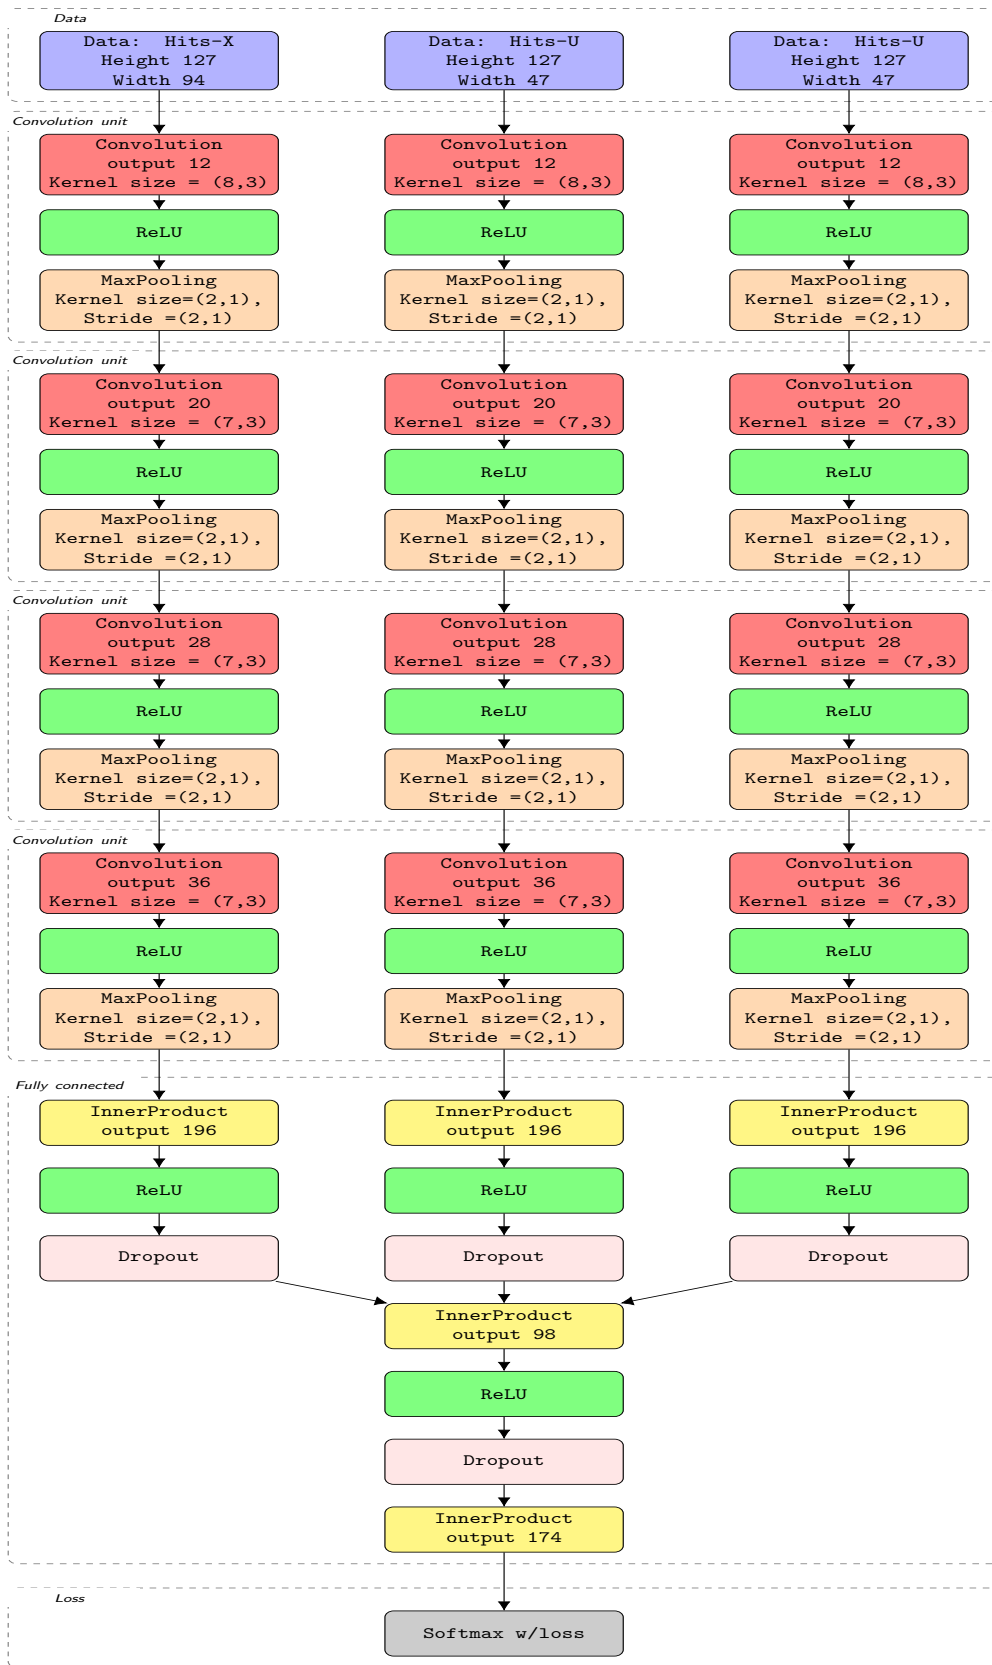


Figure 3: The flow chart of the network structure of the domain-expert-designed models.

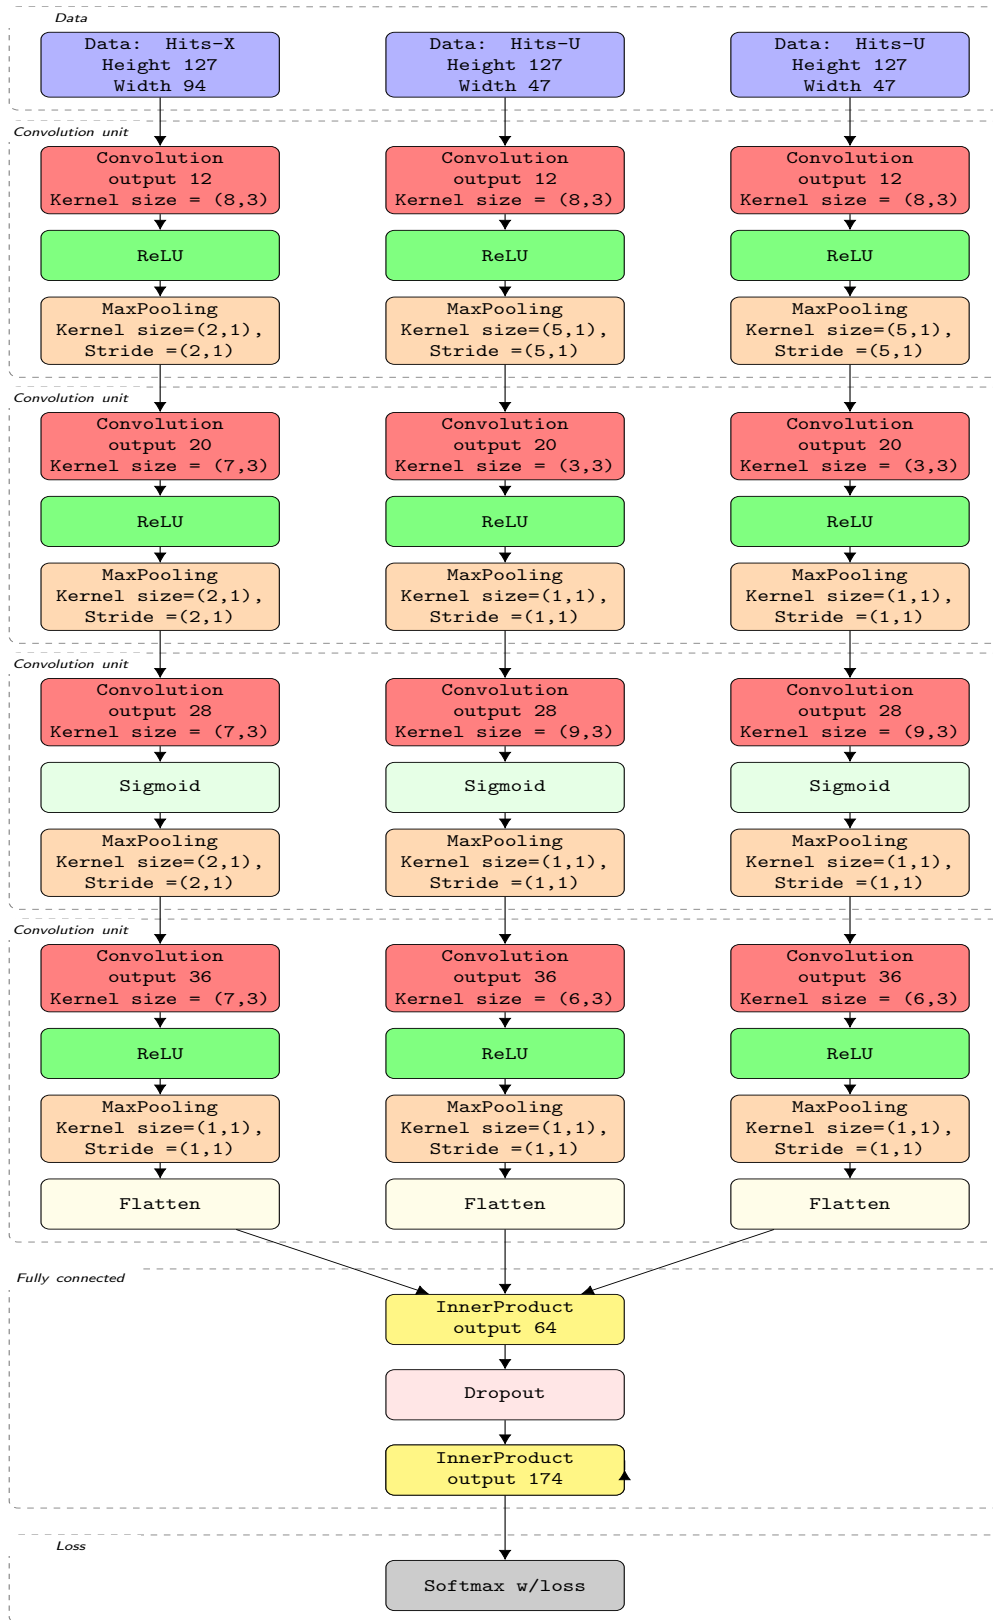


Figure 4: The flow chart of the network structure of the MENNDL model (MENNDL 68525 indicated by yellow line in Fig. 5). This model performed best among all the MENNDL models. The performance of this model is compared with the artisanal models.

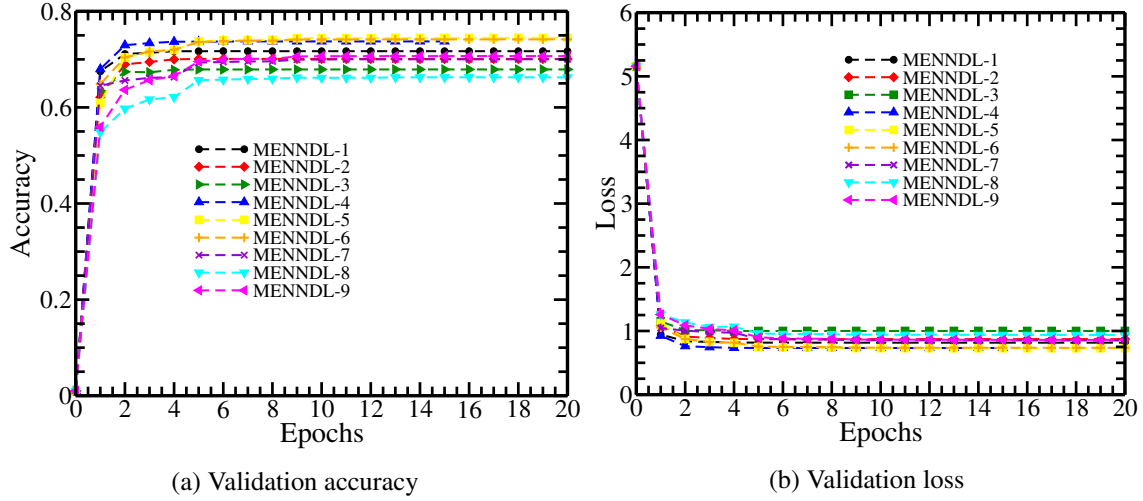


Figure 5: Validation accuracy and validation loss for different MENNDL models. Among these nine models, model *MENNDL-5* (yellow coloured square-dashed line) has the highest accuracy (74.5%) and lowest loss (0.73) values.






Figure 6: Validation accuracy and validation loss plots for domain-expert-designed(DED)-Caffe model (solid line) and MENNDL-5 model (dashed line).



(a) Domain-expert-designed(DED)-Caffe



(b) MENNDL

Figure 7: The event distributions for two different machine learning models using target 2 of the MINERvA detector. Three nuclear targets (iron, lead and carbon) from the MINERvA detector were used in this study. Events reconstructed from the source nuclear target (iron) are represented in red and other targets, lead and carbon combined, are presented in purple. Distributions of the number of events reconstructed from lead and carbon are shown by blue and indigo colors, respectively. The number of events coming from the plastic scintillator upstream (in this case six planes between targets 1 and 2, closest to target 2) and downstream (in this case six planes between targets 2 and 3, closest to target 2) of the nuclear target are also presented in orange and magenta, respectively. The bin width is 4.2 cm instead 1.7 cm but the bins are normalised to 1.7cm which is the thickness of each plane

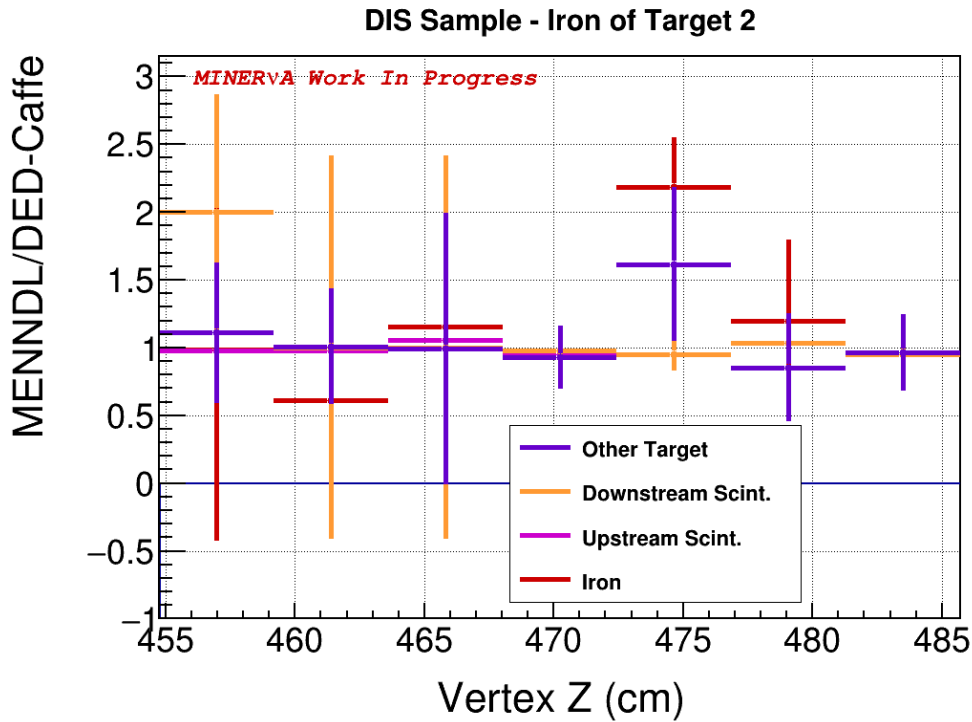


Figure 8: Ratio of the sub-figures from Fig.7.

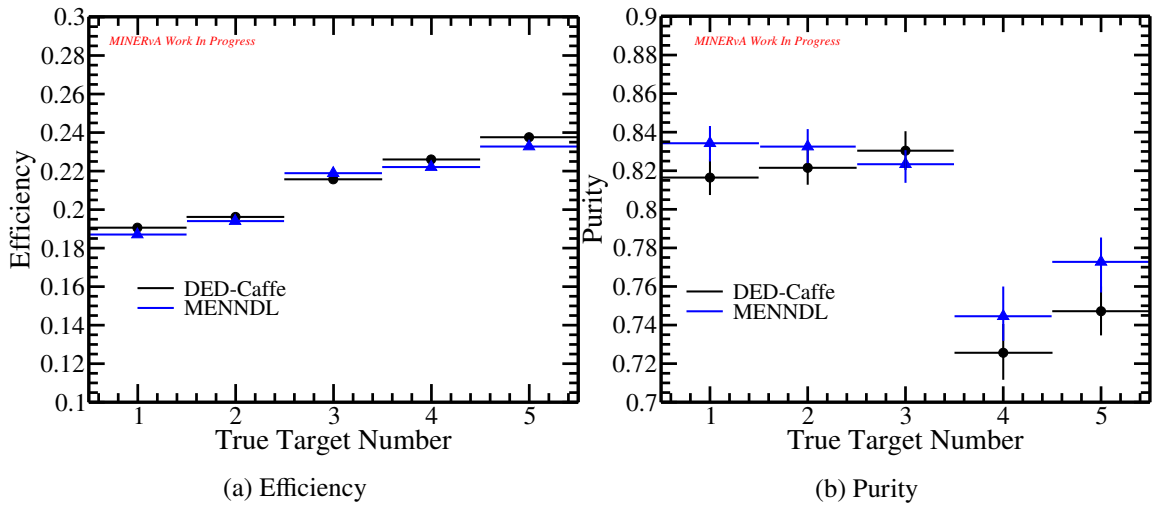


Figure 9: The results for efficiency and purity with respect to five different nuclear targets in MINERvA detector. Two different machine learning models have been used to make efficiency plots for DIS sample: domain-expert-designed(DED)-Caffe (black/circle) and MENNDL (blue/triangle).