

PAPER • OPEN ACCESS

Classification of electron and muon neutrino events for the ESSvSB near water Cherenkov detector using Graph Neural Networks

To cite this article: J. Aguilar *et al* 2025 *JINST* **20** P08030

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

You may also like

- [Multi-messenger Observations of a Binary Neutron Star Merger](#)
B. P. Abbott, R. Abbott, T. D. Abbott et al.
- [A compact Scheimpflug lidar imaging instrument for industrial diagnostics of flames](#)
Armand Dominguez, Jesper Borggren, Can Xu et al.
- [Nuclear astrophysics with radioactive ions at FAIR](#)
R Reifarh, S Altstadt, K Göbel et al.



The Electrochemical Society
Advancing solid state & electrochemical science & technology



**249th
ECS Meeting**
May 24-28, 2026
Seattle, WA, US
*Washington State
Convention Center*

Spotlight Your Science

***Submission deadline:
December 5, 2025***

SUBMIT YOUR ABSTRACT

RECEIVED: March 31, 2025

REVISED: July 10, 2025

ACCEPTED: July 23, 2025

PUBLISHED: August 26, 2025

Classification of electron and muon neutrino events for the ESS ν SB near water Cherenkov detector using Graph Neural Networks

J. Aguilar et al.

Full author list at the end of the paper

E-mail: kaare.iversen@fysik.lu.se

ABSTRACT. In the effort to obtain a precise measurement of leptonic CP-violation with the ESS ν SB experiment, accurate and fast reconstruction of detector events plays a pivotal role. In this work, we examine the possibility of replacing the currently proposed likelihood-based reconstruction method with an approach based on Graph Neural Networks (GNNs). As the likelihood-based reconstruction method is reasonably accurate but computationally expensive, one of the benefits of a Machine Learning (ML) based method is enabling fast event reconstruction in the detector development phase, allowing for easier investigation of the effects of changes to the detector design. Focusing on classification of flavour and interaction type in muon and electron events and muon- and electron neutrino interaction events, we demonstrate that the GNN reconstructs events with greater accuracy than the likelihood method for events with greater complexity, and with increased speed for all types of events. The GNN flavour classification of neutrino interaction events results in a true positive rate of 85.87 % (57.90 %) for muon (electron) neutrinos, compared to 35.55 % (0.21 %) for the likelihood-based method with identical constraints on the false positive rate, while the reconstruction speed is increased by a factor of 10^4 . Additionally, we investigate the key factors impacting reconstruction performance, and demonstrate how separation of events by pion production using another GNN classifier can benefit flavour classification.

KEYWORDS: Analysis and statistical methods; Cherenkov detectors; Neutrino detectors; Performance of High Energy Physics Detectors

ARXIV EPRINT: [2503.15247](https://arxiv.org/abs/2503.15247)



Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | Methods | 2 |
| 2.1 | Simulation | 2 |
| 2.2 | Reconstruction | 4 |
| 3 | Results | 8 |
| 3.1 | Charged lepton events | 8 |
| 3.2 | Charged current neutrino events | 10 |
| 3.3 | Neutral current neutrino events | 20 |
| 3.4 | Reconstruction speed | 21 |
| 4 | Conclusion | 22 |
| A | Events displays of events with superior GNN Performance | 23 |
| | Author List | 26 |

1 Introduction

The existence of CP violation in the leptonic sector is among the most interesting open questions in particle physics, and can be investigated through the study of neutrino oscillations [1]. The potential impact of the CP violating phase δ_{CP} on the neutrino oscillation probability is almost 3 times larger at the second oscillation maximum than at the first [2], but requires 9 times more statistics, as the second maximum is located 3 times further away than the first, and the flux is inversely proportional to the square of the distance. This can be achieved by using a more intense neutrino beam (about one order of magnitude), measuring over a longer period (about 9 times longer), or using more efficient event selection algorithms. The European Spallation Source (ESS) linear accelerator (linac) in Lund, Sweden [3] can be modified to provide such a high intensity neutrino beam, and the ESS ν SB is the experiment proposed to measure CP violation using this beam [4, 5]. ESS ν SB will measure the flavour composition of the unoscillated neutrino beam immediately after the beam target at the near detector, which consists of a water Cherenkov (WC) detector, a magnetized scintillator detector called the Super Fine-grained Detector (SFGD), and an emulsion-based detector similar to the NINJA detector employed at J-PARC [6]. The WC detector, which this study concerns, will consist of a horizontally oriented detector tank with a length of 10.8 m, a diameter 9.4 m, and a volume of 750 m³ submerged in clean water, with the cylinder axis aligned with the neutrino beam direction. Its interior will be instrumented with just over 22,000 3.5 inch diameter photo-multiplier tubes (PMTs) resulting in a $\sim 30\%$ instrumentation coverage. All the components of the near detector are illustrated in figure 1. The flavour composition of the oscillated neutrino beam will be measured at the far detector, composed of two Cherenkov tanks located 340 km from the target such that the flux will contain both the first and second maximum of muon (anti)neutrino to electron (anti)neutrino oscillations [4]. By measuring the flavour composition of both the neutrino and antineutrino beam the oscillation probability can be calculated and δ_{CP} can be estimated.




Figure 1. The three detector components of the near detector complex, the emulsion detector, the SFSGD and the water Cherenkov detector. The direction of the neutrino beam is left to right, perpendicularly to the face of the WC tank. Reproduced from [4]. CC BY 4.0.

To enable these measurements, accurate reconstruction of the neutrino events is a necessity. While accurate reconstruction methods exist and are implemented in neutrino experiments [7, 8], they are often based on likelihood methods and have limits to their performance and reconstruction speed. The latter becomes particularly important in the development stages of the ESS ν SB, as having fast event reconstruction methods enable faster exploration of different detector designs and PMT layouts. Having multiple reconstruction methods will allow for more thorough studies of robustness to noise and detector properties, and a fast reconstruction method will allow these studies to be carried out at a higher rate. Recently, Machine Learning (ML) methods and particularly methods based on Graph Neural Networks (GNNs) have proven useful for neutrino event reconstruction [9]. In the following sections, we explore the possibilities of replacing the current likelihood-based method with one based on a GNN for classification of flavour and interaction type.

2 Methods

2.1 Simulation

The GNN models described in the following sections were trained and subsequently benchmarked against the likelihood-based method applied to two data sets of Monte Carlo (MC) simulated events (used in the ESS ν SB conceptual design report (CDR) [4] and [10]). The events were simulated using the neutrino interaction vertex generator GENIE v3.0.6 [11, 12] for neutrino interactions, and the water Cherenkov simulation software WCSim [13] for particle transport, and the detector response. WCSim is based on Geant4 [14] and developed by the Hyper-Kamiokande collaboration. For these simulations the Near Detector tank is configured to have a volume of 750 m³ with a photomultiplier tube (PMT) coverage of 30%.




Figure 2. The expected neutrino flux 100 km downstream of the target station, in the central $10\text{ m} \times 10\text{ m}$ region, for neutrino- and antineutrino-mode, per running year, as functions of neutrino energy. Reproduced from [4]. CC BY 4.0.




Figure 3. Example event displays of the simulated electron (a), muon (b), electron neutrino (c), and muon neutrino (d) events.

The first dataset contains samples of single charged leptons. These events have been generated homogeneously distributed within the detector tank with an isotropic direction distribution, and uniformly covering kinetic energies from 0 GeV to 1.5 GeV. The energy range was chosen to approximately match the expected energy range of the neutrino beam, which is limited by the beam energy of the ESS linear accelerator. Figure 2 shows the expected neutrino flux 100 km downstream of the target station in the central $10\text{ m} \times 10\text{ m}$ region, which, following the approach of the CDR [4], is assumed to be representative of the neutrino flux at the near detector. Both leptons and anti-leptons were simulated in equal amounts. In the absence of neutrinos and neutrino interactions for this dataset, the flavour of the charged lepton was the target variable for classification.

The second dataset contains a uniform distribution of interacting neutrino energies from 0 GeV to 1.5 GeV, with pure water as target material. These events were classified by the flavour of the interacting neutrino and later by interaction type. The simulations contain only muon neutrino and electron neutrino events (muons and electrons for the charged lepton simulations), as the tau neutrinos were omitted from this study.¹ Both neutrinos and anti-neutrinos and events with charged current (CC) and neutral current (NC) interactions were simulated in equal amounts.

Example event displays from the two datasets are shown in figure 3.

¹With a beam energy of up to 2 GeV and the tau lepton mass $m_\tau \approx 1.777\text{ GeV}$, the tau neutrinos are not produced in significant amounts at the target, and will not partake in CC interaction in the detectors.

2.2 Reconstruction

2.2.1 Likelihood-based reconstruction

The baseline method used for classification is based on the fitQun algorithm [8, 15]. A log-likelihood-based reconstruction method developed for the Super-Kamiokande Experiment and based on methods developed for the MiniBooNE experiment [16], it can accurately reconstruct the flavour and kinematic parameters of the interaction products, from which information about their neutrino parents can be inferred.

$$L(X) = \prod_j^{\text{unhit}} P_j(\text{unhit}|\mathbf{x}) \prod_i^{\text{hit}} \{P_i(\text{hit}|\mathbf{x})\} f_q(q_i|\mathbf{x}) f_t(t_i|\mathbf{x}) \quad (2.1)$$

The likelihood-function evaluated in the fitQun algorithm is defined in Equation 2.1, and is expressed as a product over all PMTs, given the track parameters \mathbf{x} . $P_j(\text{unhit}|\mathbf{x})$ and $P_i(\text{hit}|\mathbf{x})$ represent the probability of track \mathbf{x} producing a hit in the PMT, and $f_q(q_i|\mathbf{x})$ and $f_t(t_i|\mathbf{x})$ represent the probability of recording the charge q_i at time t_i in the case of a hit, given the track parameters \mathbf{x} .

For a given event, the fitQun framework will provide a negative logarithmic likelihood (NLL) estimation for each relevant lepton flavour hypothesis. The model score in this work follows the approach of the CDR [4] and is taken to be the ratio between the NLL of an event being a muon (neutrino) event and the NLL of being an electron (neutrino) event.

$$\text{Model score} = \frac{NLL^{\text{fitQun}_\mu}}{NLL^{\text{fitQun}_e}} \quad (2.2)$$

2.2.2 Likelihood-based cuts

For the reconstructions in the CDR [4], for which the likelihood-based method was used, a number of cuts² are applied to the data to filter out events that are hard to reconstruct with good accuracy. These cuts are based on reconstructed variables from the fitQun-based reconstruction and include, as examples, the cuts shown in figure 4. On the left, the ratio between reconstructed energies using a muon (neutrino) and electron (neutrino) hypothesis, respectively, is used to filter out muon (neutrino) events where the produced muons have energy below the Cherenkov threshold. These events are primarily detected via the electron from the muon decay, and thus frequently misidentified as electron (neutrino) events. On the right, the distance from the starting initial position of the charged lepton (which would be the interaction vertex) and the detector wall along the direction of travel of the lepton, is used to filter out events where the charged lepton is produced too close to the wall, which are often misidentified because the lepton does not always lose all its energy inside the water volume.

As these cuts are based on variables from the fitQun-based reconstruction, replacing the fitQun-based reconstruction method with a GNN-based method also means losing the ability to perform these cuts. Additionally, about half of the events are lost when applying these cuts, so removing or loosening the cuts could provide better statistics. Thus, part of this work is also investigating if the GNN classification model can perform well without these likelihood-based cuts applied.

²For a full list of cuts, see [4], section 7.2.3.

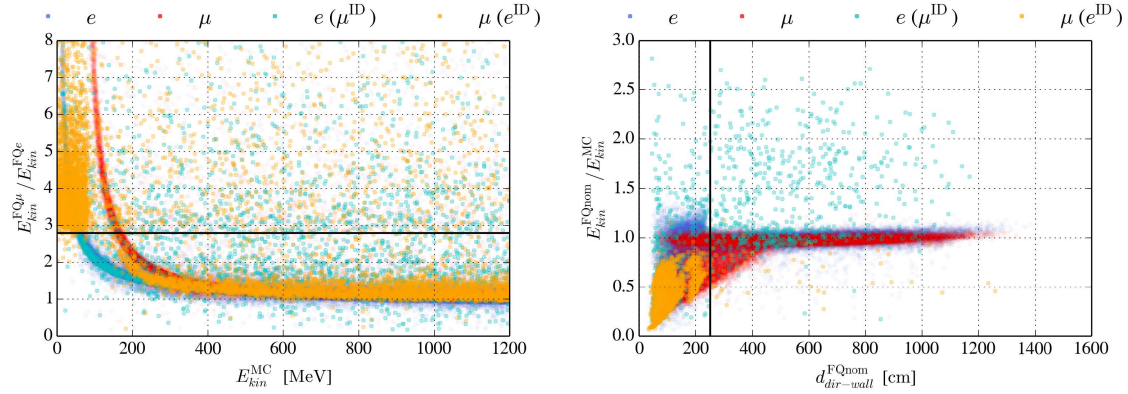


Figure 4. Illustration of two of the likelihood-based cuts used to remove incorrectly identified events from the charged lepton data, showing correctly identified electrons (dark blue) and muons (red) and incorrectly identified electrons ($e(\mu^{ID})$, cyan) and muons ($\mu(e^{ID})$, orange) distributed by reconstructed variables (and in one case, simulated energy). The classification uses the fitQun-based approach described in [4]. The sub-Cherenkov criterion (left) shows events distributed by simulated energy (E_{kin}^{MC}) and ratio between reconstructed energies using a muon ($E_{kin}^{FQ\mu}$) and electron (E_{kin}^{FQe}) hypothesis, respectively. The black line indicates the cut, (removing events with $E_{kin}^{FQ\mu} / E_{kin}^{FQe} > 2.8$), which serves to reject muon events where the muon has too little energy to produce Cherenkov radiation, and as a result is primarily detected through its electron decay product and misidentified as an electron. The Cherenkov-ring resolution criterion (right) show events distributed by distance from the interaction vertex to the detector wall along the direction of the charged lepton d_{wall}^{FQnom} , and the ratio of the reconstructed and simulated energies. The black line indicates the cut, removing events with $d_{wall}^{FQnom} < 250$ cm which are often misidentified. Reproduced from [4]. CC BY 4.0.

2.2.3 Graph Neural Networks

Various ML reconstruction methods have been explored in the context of neutrino experiments and Cherenkov detectors. Convolutional neural networks (CNNs) have been employed with positive results [17], and GNNs have seen increased use in recent years in neutrino detector design and event reconstruction. [9, 18]. GNNs are an ML method based on graph theory, in which data is expressed as graphs of nodes connected by edges, and node information is propagated between the layers of the network through the exchange of information with neighbouring nodes (known as message passing) [19]. CNNs and GNNs are similar in they way they deal with the exchange of information between data points that are related (often by proximity in a chosen feature space). Although GNN convolutions and message passing is computationally expensive — as graphs can be very large (the graphs in this study have $O(10^3)$ nodes) and computations are carried out for each pair of connected nodes — they often make up for it by having fewer layers than a typical deep CNN. Additionally, CNNs require the input to have a grid-like structure with even point-spacing, which is not typically the case for Cherenkov detectors. Adapting data structure for the use of CNNs may require padding or other modifications that leads to unnecessary computations or misrepresentation of the relationship between data points. With GNNs the spatial distances and relations between points are incorporated in the message passing, and for this work they have been deemed a good fit. In the case of this study, a neutrino interaction event (or an event containing a single charged lepton) is described by a single graph, with the nodes of the graph being the PMT signals associated with the event, with the node

Table 1. Scaling factors for the five PMT variables which act as nodes features in the GNN graphs. Each factor is chosen according to the distribution of the variables, such that the mean of the absolute values is ~ 1 , to work optimally with the GNN activation functions.

| Variable | x, y (cm) | z (cm) | Time (ns) | Charge (PE) |
|----------------|-----------------|-----------------|-----------------|---------------|
| Scaling factor | $\frac{1}{300}$ | $\frac{1}{350}$ | $\frac{1}{100}$ | $\frac{1}{5}$ |

features PMT position (x, y, z) , signal arrival time and recorded charge. The graph is completed by connecting the nodes with edges according to a K -nearest neighbours scheme (in x, y, z, t). As a result, the model input for an event with N PMT signals will be the input feature matrix $\mathbf{X} \in \mathbb{R}^{N \times 5}$, accompanied by the edge index matrix $\mathbf{E} \in \mathbb{N}^{2 \times M}$, where the number of edges M is $K \cdot N$. The input features are transformed according to a simple scheme where each feature is scaled by a constant value. The scaling factors are shown in table 1.

To train and apply GNN models to the graphs, the GraphNeT [20] framework is used. Initially developed by the IceCube collaboration [21, 22], GraphNeT facilitates training of ML models (including GNNs) on data structures and tasks commonly associated with Cherenkov detectors and neutrino telescopes, using the PyTorch [23], PyTorch Geometric [24] and Pytorch Lightning [25], libraries as backend. GraphNeT takes as input tables of PMT signals and event labels, and handles parallel loading of data and building of graphs during training, outputting trained models and predictions for the desired data sets. In addition to useful loss functions and other optimization utilities, GraphNeT contains various model architectures like the DynEdge model. The DynEdge architecture is based on the EdgeConv [26] message passing scheme, but with the addition of edge reconfiguration between each graph convolution [9], and is employed for this study without modifications. A diagram of the DynEdge architecture is shown in figure 5. The learning rate was controlled by a one-cycle linear learning rate scheduler [27] with manual tuning of the maximum learning rate and early stopping to prevent overfitting. The varying number of PMT signals and corresponding graphs sizes is handled without padding using Pytorch Geometric’s mini-batching scheme. For details, see [24]. During training, the graph building is performed on CPU cores and moved to the GPU to ensure maximum occupancy.

For each event, the GNN model will output a model score between 0 and 1, representing the probability assigned by the model that the event is of the signal flavour, where a score closer to 1 indicates higher probability of signal flavour. The selection threshold can be adjusted within the this range to obtain a satisfactory sample. For this study, the default signal flavour will be electron flavour, such that the truth value for flavour is 1 when the event is an electron neutrino event. The following sections will include plots where the muon flavour is treated as truth (such as 7, right). These are the same predictions as for electron flavour, but with the labels of the output reversed. The GNN flavour classification model developed in this study was only trained on the neutrino interaction dataset, but has been tested on both the charged lepton and neutrino datasets after training as described in the following sections.

2.2.4 Training

Both data sets (neutrinos and charged leptons) were classified by neutrino or lepton flavour (and later interaction type and pion production) using the same GNN models, which were trained only on the neutrino event data set. Training separate models applied to the charged lepton events was investigated,

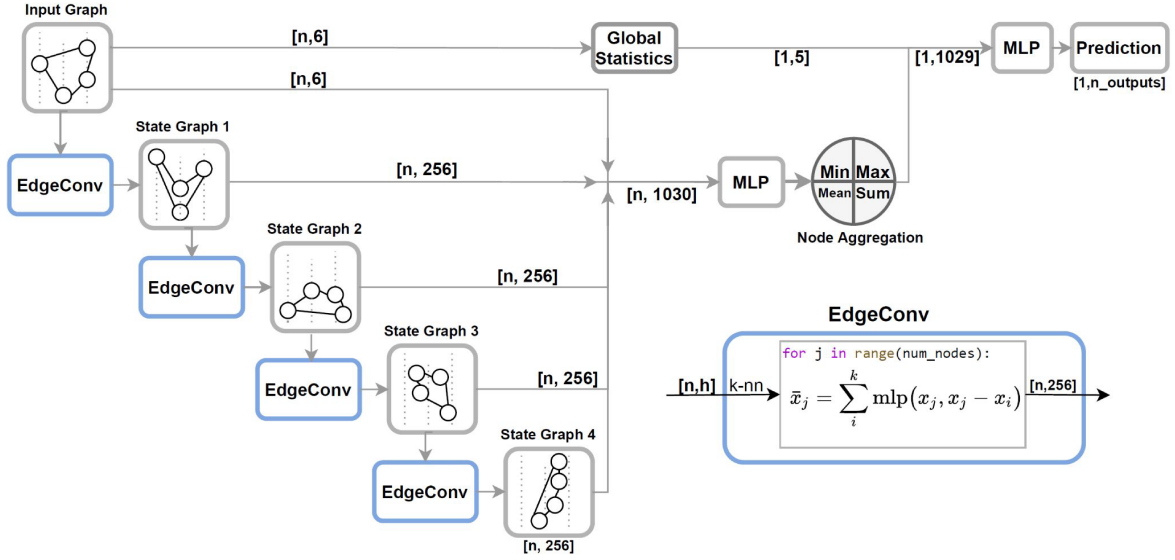


Figure 5. Diagram of the DynEdge model architecture with message passing layers on the left and aggregation layers on the right, and a detailed schematic of the EdgeConv message passing scheme. Reproduced from [9]. CC BY 4.0.

Table 2. Breakdown of the sizes of training and test samples for the charged lepton and CC neutrino interaction data sets. The charged lepton data set was not used for training, and was not tested without cuts. All samples were simulated with a 1:1 ratio of initial lepton/neutrino flavour, and a 1:1 ratio of ν and $\bar{\nu}$ for the neutrino sample. The training was divided into 67 % training data and 33 % validation data with the same ratios. For the neutrino events, a total of 800,000 events of equal parts CC and NC events were simulated.

| | Simulated | After cuts | Training | Test w/ tcuts | Test w/o cuts |
|--|----------------|----------------|----------------|----------------|----------------|
| Charged leptons | 100,000 | 49,325 | - | 49,325 | - |
| $\mu / \bar{\mu} / e / \bar{e}$ | 25,000 | - | - | - | - |
| Neutrinos (CC) | 400,000 | 226,425 | 100,000 | 100,000 | 300,000 |
| $\nu_{\mu} / \bar{\nu}_{\mu} / \nu_e / \bar{\nu}_e$ (CC) | 100,000 | - | 25,000 | - | 75,000 |
| Neutrinos (NC) | 400,000 | - | 100,000 | - | 300,000 |
| $\nu_{\mu} / \bar{\nu}_{\mu} / \nu_e / \bar{\nu}_e$ (NC) | 100,000 | - | 25,000 | - | 75,000 |

but did not yield any significant increase in performance. The models were trained on 100,000 (67 % training, 33 % validation) events on two NVIDIA Tesla K20 cards and evaluated on $\sim 50,000$ – $300,000$ events (depending on the cuts applied and the resulting sample sizes). For a more detailed overview of training and test samples, see table 2. The training was carried out using a linear 1cycle learning rate scheduler [28] with early stopping, with all models converging within 50 epochs.

2.2.5 Performance metric

The neutrino beam produced at the ESS ν SB target will contain $> 98\%$ muon neutrinos [4] with the remainder dominated by electron-neutrino. The main purpose of the Near Detector (on which this study is focused) is to accurately determine the electron neutrino content of the initial neutrino beam (before oscillations) in order to reduce systematic uncertainties for the oscillation analysis

Table 3. Confusion matrices for flavour classification with the signal flavours ν_μ (a) and ν_e (b).

| Prediction | | | | Prediction | | | |
|----------------------|-----------|---------------------|---------------------|--------------------|-----------|---------|-----------|
| | | ν_μ | ν_e | | | ν_e | ν_μ |
| Truth | ν_μ | True Positive (TP) | False Negative (FN) | Truth | ν_e | TP | FN |
| | ν_e | False Positive (FP) | True Negative (TN) | | ν_μ | FP | TN |
| (a) ν_μ signal | | | | (b) ν_e signal | | | |

conducted with far detector data. For these reasons, the purity of the electron neutrino sample (which is of greater importance) is more severely affected by the misidentification of muon neutrino events, making it more important to filter out muon neutrino events from the electron neutrino sample than electron neutrinos from the muon neutrino sample. Thus, this work follows the example of the CDR [4] and [10] and uses a false positive rate (FPR) of 0.1 % for (muons neutrinos identified as) electron neutrinos and 1 % for (electron neutrinos identified as) muon neutrinos as the target parameter when optimizing the selection criteria. Consequently, the best performing classifier will be the one that provides the highest true positive rates (TPRs) when enforcing an FPR of 0.1 % for electron neutrinos and 1 % for muon neutrinos.

The TPR is the same as the efficiency, while the FPR is the probability of labelling an event of the background flavour as signal, quantified e.g. for electron neutrinos, as the number of muon neutrino events identified as electron neutrino events relative to the total number of muon neutrinos events, and is thus independent of the ratio between the number of events of each flavour. Purity could also be used for this study, as a pure electron neutrino sample is the key objective, but since we exactly want to measure the beam composition because it is not known precisely and it is not reflected in the simulations, the TPR and FPR are better suited. For illustration, assuming $\text{TPR}_{\nu_e} = 90\%$ (and a beam composition of $f_{\nu_\mu} = 98\%$ and $f_{\nu_e} = 2\%$), the target $\text{FPR}_{\nu_\mu} = 0.1\%$ will yield an electron neutrino sample of 95 % purity, demonstrated in Equation 2.4. Extended definitions of true and false positive and negative and their dependence on the signal flavour for flavour classification are illustrated in the confusion matrices in table 3.

$$n_{\nu_\mu} = f_{\nu_\mu} \cdot \text{FPR}_{\nu_\mu} = 98\% \cdot 0.1\% = 0.098\% \quad n_{\nu_e} = f_{\nu_e} \cdot \text{TPR}_{\nu_e} = 2\% \cdot 90\% = 1.8\% \quad (2.3)$$

$$\text{purity}_{\nu_e} = \frac{n_{\nu_e}}{n_{\nu_e} + n_{\nu_\mu}} = \frac{1.8\%}{1.8\% + 0.098\%} = 95\% \quad (2.4)$$

3 Results

3.1 Charged lepton events

For charged lepton event simulations, the likelihood-based cuts described in 2.2.2 were applied to the data set to have a baseline case to compare both with the CDR results and subsequent studies in this paper. The resulting sample should have a large fraction of unambiguous events that can be easily classified by both reconstruction methods. In the case where this study proves successful and the GNN-based method can fully replace the fitQun-based method, the scenario covered in this section

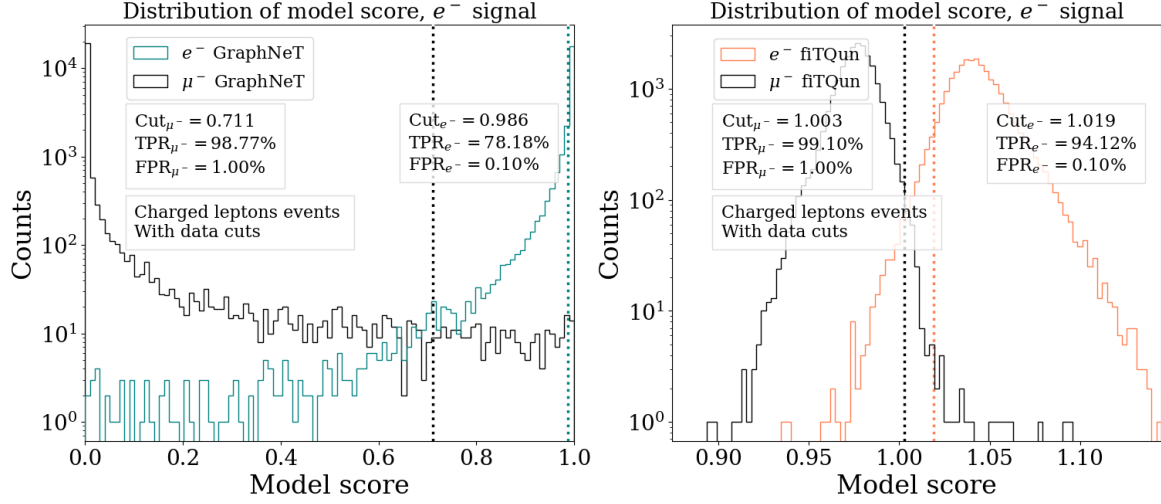


Figure 6. Model score distributions for the GNN (left) and fiTQun-based (right) lepton flavour classification method applied to the charged lepton events (with likelihood-based cuts). The GNN has been trained on neutrino events with electron flavour as signal, and the GNN model score is the network output. The fiTQun-based model score is calculated as described in 2.2.1. Dashed lines represent thresholds corresponding to target FPRs, and the resulting TPR is shown.

(and other sections where the cuts are employed) will not be possible, as we will not have access to the reconstructed variables for the cuts. These conditions serve to investigate and compare the performance of the two models under different conditions, and of the importance of the cuts themselves.

The model score distribution of the charged lepton event classification is shown in figure 6 for the GNN (left) and the fiTQun-based method (right). Vertical dashed lines indicate the threshold below (above) which events should be rejected to obtain the FPR rate described in section 2.2.5 for electrons (muons), and the resulting TPR rates are shown. Both the fiTQun-based and GNN classification methods produce acceptable results on par with the performance described in the CDR, with the fiTQun-based method having slightly better performance, yielding TPRs of 99.10 % and 94.12 % for the chosen FPR rates for muons and electrons, respectively, against the 98.77 % and 78.18 % TPRs of the GraphNeT-based method. This is to be expected, since this is the type of events the likelihood-based method is optimized for.

Figure 7 shows Receiver Operating Characteristic (ROC) curves for both reconstruction method. ROC curves are obtained by varying the model score threshold and plotting the obtained FPR (x-axis) and TPR (y-axis) and thus illustrate the trade-off between the FPR and TPR for a binary classification model. The shape of the ROC curve with the electron flavour as signal (figure 7, left) indicates that for looser constraints on FPR ($> 5 \cdot 10^{-4}$) the fiTQun-based method yields greater TPRs than the GNN, while for lower on FPRs, the GNN starts to outperform the fiTQun-based method. For muons (figure 7, right) the performance of the two methods is similar above $\text{FPR} = 10^{-2}$, while for lower target rates, the GNN has superior performance. Figure 7 also shows the calculated area under the ROC curves (AUC). The AUC describes a classification method’s ability to yield a high TPR at the full FPR range, with a perfect ROC curve being a 90° angle (or a horizontal line when the FPR-axis is logarithmic) with $\text{AUC} = 1$. The AUC is useful as a single statistic to compare the performance of reconstruction methods, but for more rigorous testing one should focus on the specific FPR requirements.




Figure 7. ROC curves for the GNN and fiTQun-based lepton flavour classification methods applied to the charged lepton events (with likelihood-based cuts) with electrons as signal (left) and muons as signal (right).




Figure 8. Confusion matrices for GNN and fiTQun-based lepton flavour classification methods applied to the charged lepton events (with likelihood-based cuts) with electrons as signal (a and b) and muons as signal (c and d).

To further illustrate the result of imposing the thresholds obtained in figure 6, confusion matrices for classification using these thresholds are shown in figure 8 for both models. Here it is also clear that large fractions of the target flavour can be accepted using both models, although more electrons are lost when using the GNN-based method (5239 events, 21.8 %) than with the fiTQun-based method (1412 events, 5.9 %), corresponding to the TPR rates obtained in figure 6.

3.2 Charged current neutrino events

3.2.1 With likelihood-based cuts

In this section the performance of the GNN and fiTQun-based classification methods applied to full neutrino interaction event simulations with the likelihood-based cuts described in 2.2.2 applied is shown, in order to illustrate how each method performs on neutrino simulations under optimal conditions.

The model score for the neutrino event classification distributed by true (simulated) neutrino energy is shown in figure 9 for the GNN (left) and the fiTQun-based method (right). The GNN-based classification yields TPRs of 96.98 % and 69.65 % for muon and electron neutrinos respectively, when applying the same restrictions on FPR, which is comparable to the performance of the reconstruction described in the CDR [4]. The likelihood-based classification yields a similar TPR of 92.86 % for muon neutrinos, but a significantly lower TPR of 0.53 % for electron neutrinos. Although the two

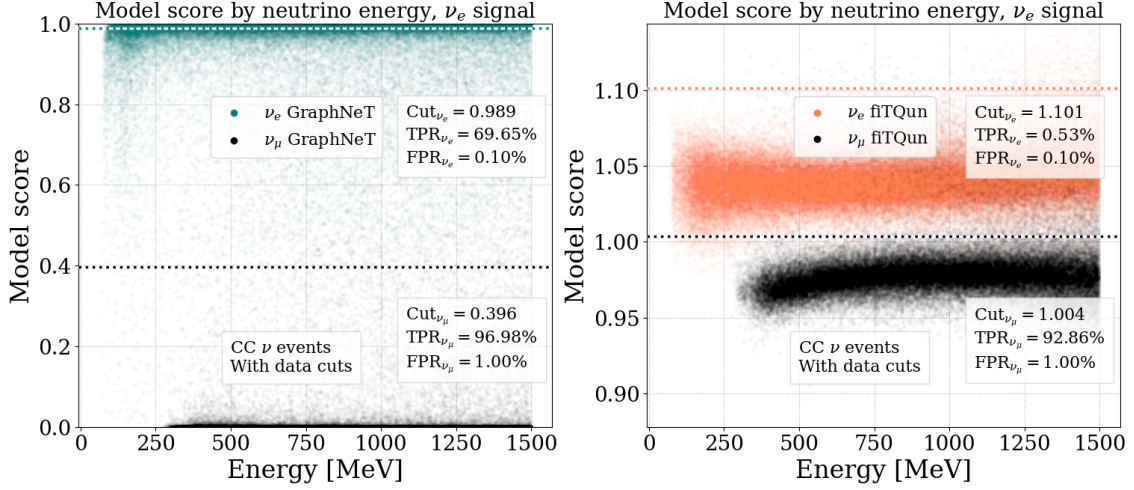


Figure 9. Model score distributed by true (simulated) neutrino energy for the GNN (left) and fitQun-based (right) neutrino flavour classification methods applied to CC neutrino events with likelihood-based cuts applied. The GNN has been trained with electron flavour as signal, and the GNN model score is the network output. The fitQun-based model score is calculated as described in 2.2.1. Dashed lines represent thresholds corresponding to target FPRs, and the resulting TPR is shown.




Figure 10. ROC curves for the GNN and fitQun-based neutrino flavour classification methods applied to CC neutrino events with likelihood-based cuts applied with electron neutrinos as signal (left) and muon neutrinos as signal (right). Similarly to the previous data sets, the ROC curves become jagged in the low FPR region, due to low statistics, but notably the ROC curve for the GNN on ν_μ becomes flat at around 0.01 % FPR, indicating the about 25 % of the muon neutrinos receive such a high model score that all the electron neutrinos can be filtered out without losing this fraction of muon neutrinos.

groups seem well separated by the likelihood-based method with no clear structure to the misidentified events, the 0.1 % FPR for electron neutrinos is too tight to produce a useful TPR.

Figure 10 shows ROC curves for both reconstruction methods, confirming that the GNN consistently outperforms the fitQun-based method for all relevant FPR rates, and that the performance difference is the greatest for lower (10^{-2} – 10^{-4}) FPR ranges of electron neutrino events. This suggests




Figure 11. Confusion matrices for GNN and fiTQun-based neutrino flavour classification methods applied to CC neutrino events with likelihood-based cuts applied with electron neutrinos as signal (a and b) and muon neutrinos as signal (c and d).

that there is a large group of muon neutrino events that are falsely identified as electron neutrinos by fiTQun, but that the GNN can correctly identify. We return to this point in section 3.2.3. Figure 11 shows confusion matrices for the two classification methods using the thresholds obtained in figure 9, showing that for the GNN-based classification with muon neutrinos as signal (c), most of the signal is kept (126,715 events, 85.9 %). For the GNN-based method with electron neutrino signal (a) and the fiTQun-based method with muon neutrino signal (d) roughly half of the signal is lost, while for the fiTQun-based classification with electron neutrinos as signal, almost all the signal is lost in order to reject the background sufficiently. While this indicates that the fiTQun based mode fails to separate the two classes, the ROC curves in figure 10 and the distribution in figure 9 show that decreasing the constraints on electron neutrino FPR would allow most of the events to be kept.

3.2.2 Without likelihood-based cuts

In the following, the likelihood-based cuts described in 2.2.2 were disregarded to investigate if the GNNs can be applied to the data without these cuts, which as argued before is important if the GNNs are to be used to optimise the detector geometry. The model score distributions of the neutrino event classification is shown in figure 12 for the GNN (left) and the fiTQun-based method (right). For this sample, the performance of both methods is similar to the performance with the likelihood-based cuts applied, though it is worse in all cases. The GNN classification yields TPRs of 85.87 % and 57.90 % for muon and electron neutrinos respectively, and the likelihood-based classification yields lower TPRs of 35.55 % and 0.21 % for muon and electron neutrinos respectively.

The model score distributed by true (simulated) neutrino energy shown in figures 13 for the GNN (left) and the fiTQun-based method (right), shows a clear band of muon neutrinos in the lowest energy range, that overlaps completely with the electron neutrinos, making it hard to exclude them without losing the majority of electron neutrinos. Since muons produced in interactions in this range will have low enough energy that the dominant part of their event signature will be from the electrons produced from their decay, such a band is expected. While the GNN performance also appears to be worst in the low energy range, such a band is not present. Instead, many of the ambiguous events obtain a model score of about 0.4, leading to them being excluded when applying the restrictions on FPR. This indicates that the GNN, from being exposed to enough of these events during training, can learn to identify them as ambiguous and assign them a model score that represents this.

Figure 14 shows ROC curves for both methods, confirming again that the GNN consistently outperforms the fiTQun-based method, and that while the difference in performance is still greatest for

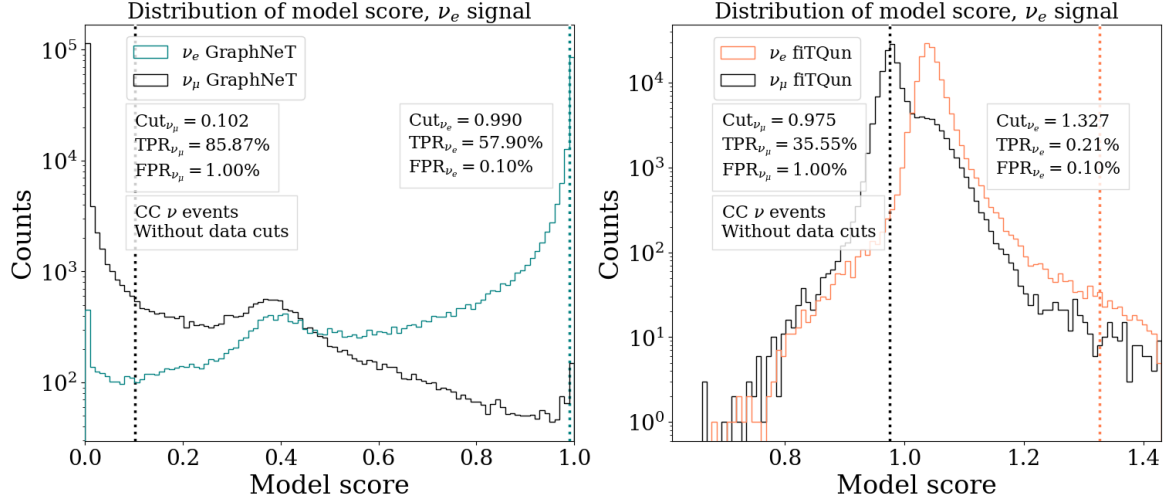


Figure 12. Model score distributions for the GNN (left) and fitQun-based (right) neutrino flavour classification methods applied to CC neutrino events without likelihood-based cuts applied. The GNN has been trained with electron flavour as signal, and the GNN model score is the network output. The fitQun-based model score is calculated as described in 2.2.1. Dashed lines represent thresholds corresponding to target FPRs, and the resulting TPR is shown.




Figure 13. Model score distributed by true (simulated) neutrino energy for the GNN (left) and fitQun-based (right) neutrino flavour classification methods applied to CC neutrino events without likelihood-based cuts applied. The GNN has been trained with electron flavour as signal, and the GNN model score is the network output. The fitQun-based model score is calculated as described in 2.2.1. Dashed lines represent thresholds corresponding to target FPRs, and the resulting TPR is shown.

lower FPR electron neutrino events, the performance difference for muon neutrino events is greater than with the likelihood-based cuts applied. Figure 15 shows confusion matrices for the two classification methods using the thresholds obtained in figure 12, showing that for the GNN-based classification with muon neutrinos as signal (c), most of the signal is kept (126,715 events, 85.9 %). For the GNN-based method with electron neutrino signal (a) and the fitQun-based method with muon neutrino signal (d)

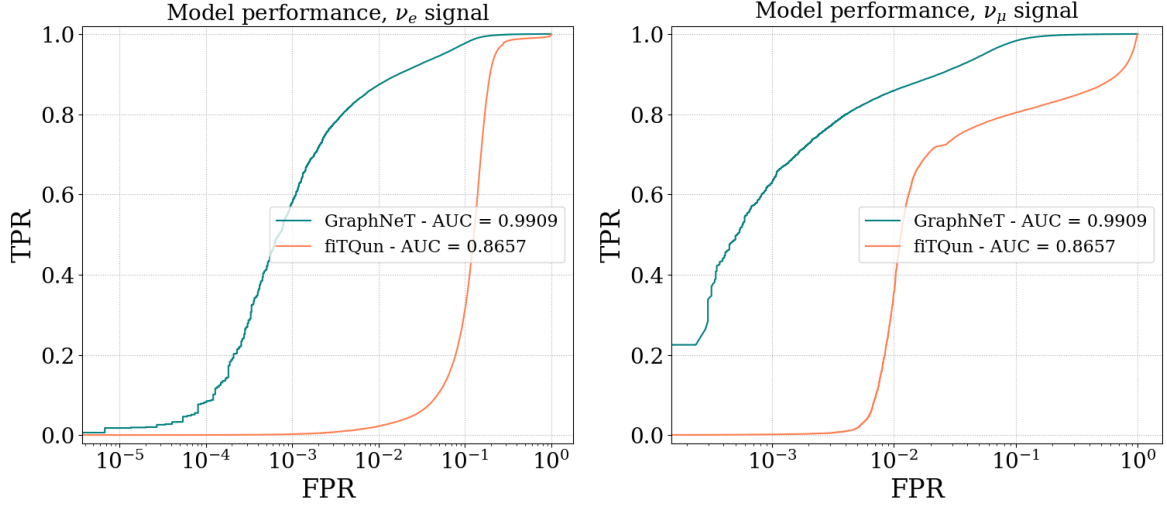


Figure 14. ROC curves for the GNN and fitQun-based neutrino flavour classification methods applied to CC neutrino events without likelihood-based cuts applied with electron neutrinos as signal (left) and muon neutrinos as signal (right).




Figure 15. Confusion matrices for GNN and fitQun-based neutrino flavour classification methods applied to CC neutrino events without likelihood-based cuts applied with electron neutrinos as signal (a and b) and muon neutrinos as signal (c and d).

roughly half of the signal is lost, while for the fitQun-based classification with electron neutrinos as signal, almost all the signal is lost in order to reject the background sufficiently.

In table 4 is shown the number of events of each flavour that are classified as either ν_e or ν_μ by both methods, the GNN-based but not the fitQun-based method, or the fitQun-based but not the fitQun-based method using the thresholds from figure 13. Confirming the previous findings, a large number of muon neutrinos are correctly identified by both methods, while the false positives are low due to the chosen thresholds. Additionally, the GNN-based method is able to correctly classify significant samples of electron neutrinos and muon neutrinos that are not accepted using the fitQun-based method.

3.2.3 Performance comparison

Many factors impact the performance of the reconstruction methods applied to the complex neutrino interaction events. In the following, a few examples of events where the two methods have similar or dissimilar performance are inspected. While the examples do not exhaust all the relevant event topologies, they can provide valuable insight about features to consider when designing the ESS ν SB reconstruction pipeline. Figure 16 shows event displays in 3D (top) and 2D (bottom) of selected

Table 4. Overview of the number of CC neutrino events without likelihood-based cuts of each flavour (ν_e and ν_μ) that are classified as either ν_e or ν_μ by both methods, the GNN-based but not the fitQun-based method, or the fitQun-based but not the GNN-based method, when using the thresholds obtained in figure 13.

| | | GNN & fitQun | GNN & NOT fitQun | fitQun & NOT GNN |
|----------------|-------------------|--------------|------------------|------------------|
| True ν_e | ν_e ID (TP) | 191 | 85,372 | 126 |
| | ν_μ ID (FP) | 102 | 1,375 | 1,375 |
| True ν_μ | ν_μ ID (TP) | 50,766 | 75,949 | 1,695 |
| | ν_e ID (FP) | 0 | 147 | 147 |

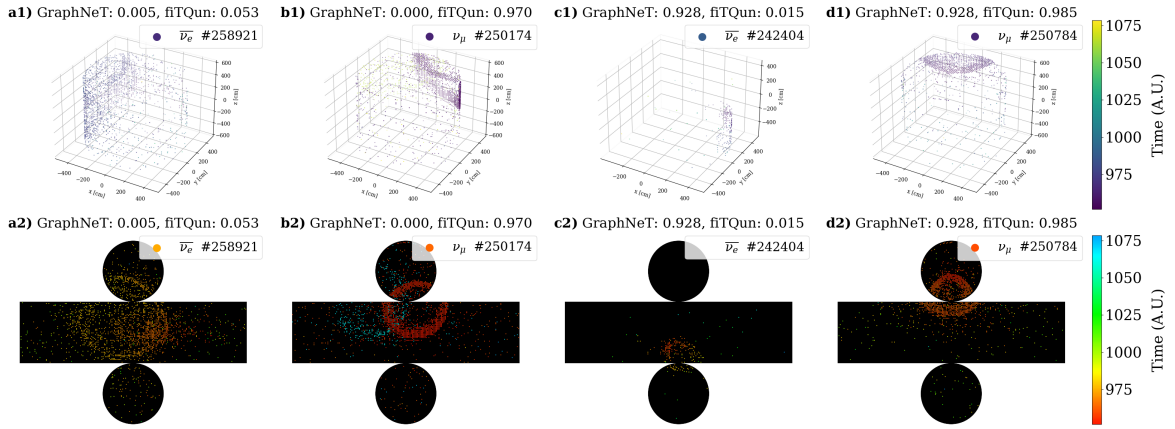


Figure 16. 3D (top) and 2D (bottom) event displays of events with different combinations of Confidence-Correctness Scores for each neutrino flavour classification method as explained in the text. The event displays show PMT signals distributed by the spatial (x, y, z) position of the PMT, with the colour indicating the arrival time of the signal.

neutrino events, created to investigate the difference in performance on neutrino data between the two methods. The events were selected by the difference between the model score and the truth label (0 for ν_μ , 1 for ν_e), here denoted the Confidence-Correctness Score (CCS), for each method. The CCS, illustrated in Equation 3.1, is a measure of how “mistaken” the classification method is for a given event, regardless of the flavour of the event. A CCS of ≈ 0 (a correctly confident classification) can either be a ν_μ event (0) with a model score of ≈ 0 , or a ν_e event (1) with a model score of ≈ 1 . In contrast, a CCS of ≈ 1 signifies a confidently incorrect prediction, and arises from a model score of ≈ 0 for a ν_e event (1) or vice versa. In the following, the events were selected to have different combinations of CCS’s for the two reconstruction methods.³

$$\text{CCS} = |\text{truth} - \text{prediction}| \quad (3.1)$$

³Combinations of: low-low, low-high, high-low, high-high, where a low score (CCS < 0.1) indicates a correct prediction with high confidence, and a high score (CCS > 0.9) indicates an incorrect prediction with high confidence. The value 0.1 was chosen to accept most of the unambiguous events and the value 0.9 = 1 – 0.1 is derived from that.

The first event from the left (figure 16, a1 and a2) shows an event where both methods predict correctly with great confidence. This event has many detector hits, providing lots of information to the reconstruction algorithms, as well as a clear ring structure. The second event (figure 16, b1 and b2), where the GNN-based method predicts correctly and the fitQun-based method incorrectly, is of particular interest, as both the 2D and 3D event displays show two distinct Cherenkov Rings with a significant time delay, most likely from a event where a muon neutrino produces a muon that decays and produces an electron. This indicates that the GNN has learned to identify these more complex events out of the box, even for cases where it is difficult for fitQun, a finding that is consistent when looking at figure 25 of appendix A, which shows several event displays of similar CCS combination, many of which exhibit double rings. The third event (figure 16, c1 and c2) where the fitQun-based method predicts correctly with great confidence, and the GNN predicts incorrectly, has few detector hits, but a clear ring structure. This can be interpreted as the event not having enough data points for the GNN - which as an ML model needs more information to function optimally - to identify the pattern, but ring-like enough for the fitQun-based method - which as an LLH-based method does not necessarily have the same limitations - to converge. The final event (figure 16, d1 and d2) has an average amount of detector hits, and a clear ring shape, and should be easily interpretable by both methods. It is not clear why both methods have poor performance on such an event, and future studies should look into this behaviour. Since the main focus of this section is on the events where the methods differ in performance, it is not investigated further here.

3.2.4 Pion production

Another factor that blurs the picture of neutrino interactions in the ESS ν SB detector is the production of pions. Pions are a common product of weak neutrino-nucleus interactions, created both in resonance production processes and coherent scattering. Both CC and NC coherent scattering can produce pions, with CC processing being the dominant contributor. The pion production can interfere with measurements, both because they can create a signature that looks like Cherenkov radiation from a charged lepton, and because their energy will be missing from the charged lepton energy reconstruction and consequently from the neutrino energy estimation [29, 30].

Both events with and without pion production are present in the neutrino event dataset, and from the simulation output, the number of pions created can be obtained. Figure 17 shows the ROC curves for the GNN and fitQun-based methods applied to the neutrino events (in this case with the likelihood-based cuts, to better visualise the differences in performance between the two methods), but separated into events where pions are produced in the final state (which make up $\sim 10\%$ of the events) and events where no pions are produced. It is clear, that both methods are able to classify events with no pion production, whereas for the event with pion production, the fitQun-based method performs significantly worse than the GNN for both electron and muon neutrinos. This is also demonstrated in figure 18 which show histograms of the model scores for the GNN (left) and fitQun-based (right) classification methods. For the GNN, there is still a good separation of the events with pions, but the number of events receiving a score near 0 or near 1 is lower, while the number of events receiving scores in the rest of the spectrum is higher, signifying that the events are harder to classify. For the fitQun-based method, the distribution of ν_e events is unchanged, but the distribution of ν_μ events are shifted to the right (higher model scores), resulting in a large overlap of ν_e and ν_μ events. This indicates, that one might want to start by using classification model (like another GNN) to separate




Figure 17. ROC curves for the GNN and fitQun-based neutrino flavour classification methods applied to CC neutrino events (with likelihood-based cuts) with electron neutrinos as signal (left) and muon neutrinos as signal (right) in the cases where pions and no pions are produced in the final state.

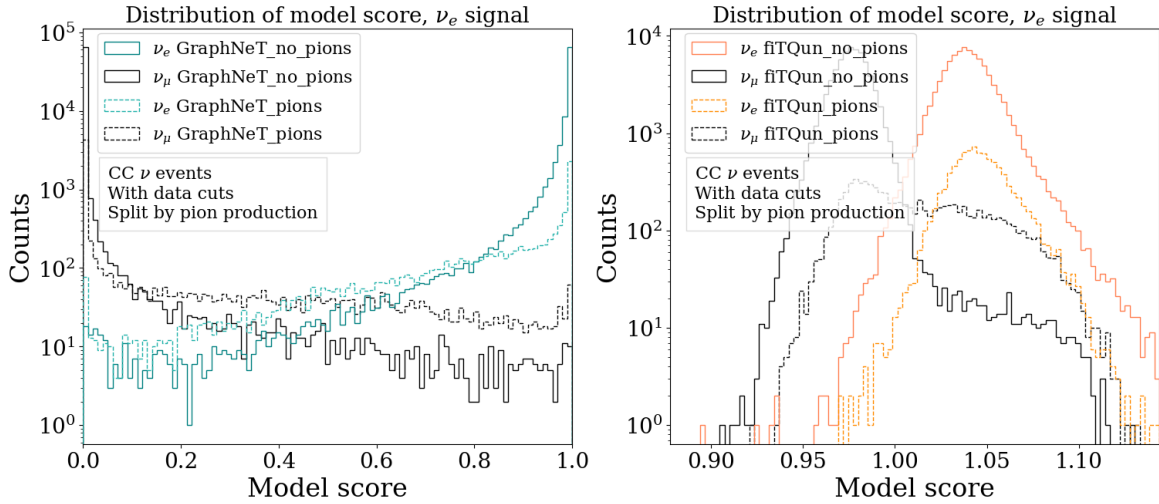


Figure 18. Model score distributions for the GNN (left) and fitQun-based (right) neutrino flavour classification methods applied to CC neutrino events (with likelihood-based cuts) in the cases where pions and no pions are produced in the final state. The GNN has been trained with electron flavour as signal, and the GNN model score is the network output. The fitQun-based model score is calculated as described in 2.2.1.

events into classes with or without pions before running the reconstruction, in order to treat the two cases separately. We will test this idea in the next subsection.

3.2.5 Pion classifier

Following the results in section 3.2.4, a GNN classification model was trained on the CC (CC) neutrino interaction data set to distinguish between events with and without pion production. The model target variable was chosen to be a boolean that had the value true (1) for events where any number of pions were created in the interaction, and false (0) for events where no pions were produced. Neutrino

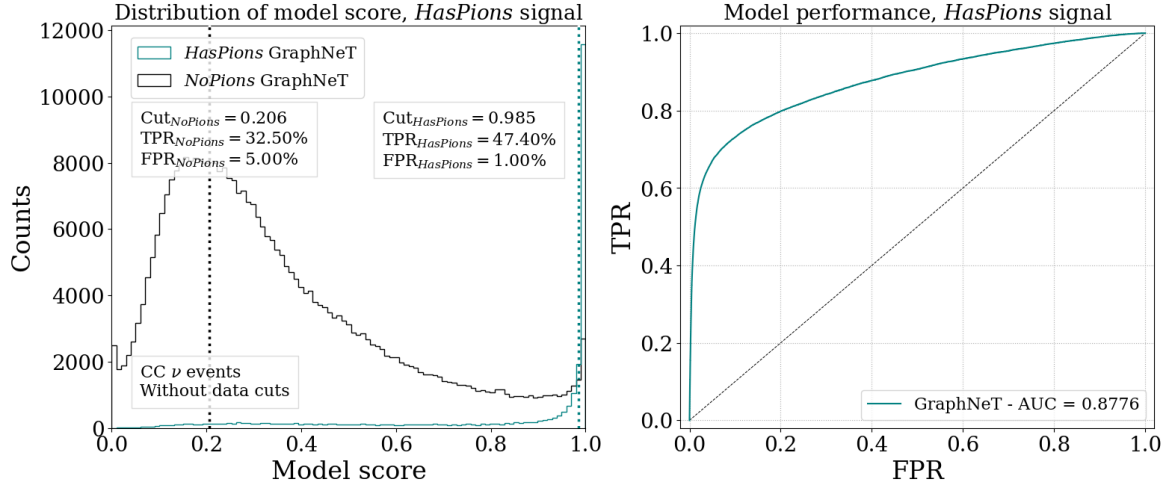


Figure 19. Model score distribution (left) and ROC curve (right) for the GNN pion production classification methods applied to CC neutrino events. The GNN has been trained with one or more pions in the final state as signal, and the GNN model score is the network output. The fitQun-based model score is calculated similarly to the calculation described in 2.2.1. Dashed lines on the left represent cuts corresponding to chosen FPRs, and the resulting TPR is shown.

flavour was not taken into account for this model, and the likelihood-based cuts were not applied, as the goal was to obtain the most general model possible to separate all types of events by pion production. The model score distribution and the ROC curve for the resulting model, are shown in figure 19 on the left and right side, respectively. Together, the figures show that while the model can classify most events correctly, it is harder for the model to confidently identify events without pion production than it is for events with pion production. This can be seen from the model score distribution on the left, as almost all the events with pion production are assigned a model score close to 1, whereas the events without pion production receive a wide range of model scores. Consequently, for a FPR requirement of 1.00 %, the model can retain 47.4 % of the sample of events with pions. For events without pion production, the majority of the sample is lost when enforcing an FPR of 1.00 %. Instead, figure 19 shows the result of raising the requirement to 5.00 % FPR, which produces allow a TPR of 32.5 %. The impact of using this classifier will be investigated in the following section.

3.2.6 Flavour classification performance split by the pion model

With the pion classifier described in the previous section completed, the next objective was to study how the reconstruction for both the GNN and the fitQun-based method was impacted by filtering the data with this classifier. For this, both the GNN and fitQun-based flavour classification methods were applied to the neutrino events (in this case with the likelihood-based cuts, to better visualise the differences in performance between the two methods), but separated by the pion production classification model described in section 3.2.5, using the thresholds shown on the left of figure 19, to obtain reasonably pure samples of events with and without pion production. Figure 20 shows the ROC curves for two models and the two data samples, which are almost identical to the ones shown in figure 17, showing that both methods perform very well on the events without pion production, and that the GNN has significantly better performance on the events with pion production than the




Figure 20. ROC curves for the GNN and fitQun-based neutrino flavour classification methods applied to CC neutrino events (with likelihood-based cuts) with electron neutrinos as signal (left) and muon neutrinos as signal (right) separated by events with and without pion production using the pion production classification model.

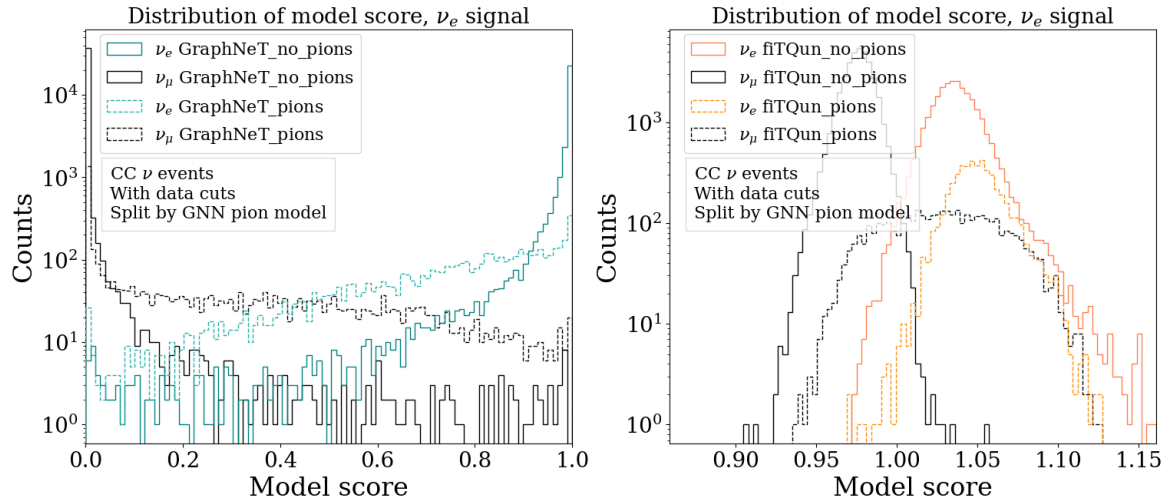


Figure 21. Model score distributions for the for the GNN (left) and fitQun-based (right) neutrino flavour classification method applied to CC neutrino events (with likelihood-based cuts) separated by events with and without pion production using the pion production classification model. The GNN has been trained with electron flavour as signal, and the GNN model score is the network output. The fitQun-based model score is calculated as described in 2.2.1.

fitQun-based method. This is confirmed when looking at figures 21 which shows the model score distributions for the GNN (left) and the fitQun-based method (right). For the GNN, while the overlap is greater for the events with pion production than for the events without, there is still a significant separation for both samples. For the fitQun-based method, the separation of events without pions is very clear, whereas there is a large overlap between the two flavours for the events with pion production. This indicates that the pion production classifier is working as intended, separating events into categories with different reconstruction difficulty.




Figure 22. ROC curves for the GNN and fiTQun-based neutrino flavour classification methods applied to CC and NC neutrino events with electron neutrinos as signal (left) and muon neutrinos as signal (right) separated by interaction type.

3.3 Neutral current neutrino events

For the next section, the neutrino interaction dataset with both CC and NC events is used without any likelihood-based cuts. This includes 400.000 simulated NC events in addition to the CC events introduced in table 2 and used in the previous sections. A GNN neutrino flavour classification model was trained on 100.000 events with equal representation of interaction types, flavours and neutrinos/antineutrinos, and tested on the remaining events. Figure 22 shows the ROC curves for this GNN and the fiTQun-based method. The data for the ROC curves is split by CC and NC interactions, and figure 22 shows clearly, that for the NC dataset both methods have poor performance comparable to a random selection. This is also demonstrated in figures 23 which shows histograms of the model scores for the GNN (left) and fiTQun-based (right) classification methods, with the same data separation. For both methods there is a complete overlap in model scores for the NC events. An interesting observation is that the GNN scores almost all NC events close to 0.5, indicating that the model is unable to learn patterns that separate the events. As the NC events contain no primary charged lepton product, their topology is dependent only on the kinematics of the event and not the flavour of the incoming neutrino. which explains why both methods fail to identify the flavour of these events. Without a reliable way to classify these events by flavour, the development of a method for filtering out the NC events and only analyse the CC events could be beneficial.

3.3.1 Separation of CC and NC events

In an effort to meet the needs described earlier in this section, a GNN model for classifying events according to the interaction type was trained on the neutrino interaction data set with both interaction types represented, and the results are shown in figure 24. As this is not a feature of the fiTQun-based method, no comparison is shown. The ROC curve in the right plot of figure 24 shows promising performance, as 99 % of NC events can be discarded while keeping a little less than half of the CC events. The histogram on the left of figure 24 shows a greater spread of the model score for the




Figure 23. Model score distribution for the GNN (left) and fiTQun-based (right) neutrino flavour classification method applied to CC and NC neutrino events separated by interaction type. The GNN has been trained with electron flavour as signal, and the GNN model score is the network output. The fiTQun-based model score is calculated as described in 2.2.1.




Figure 24. Model score distribution (left) and ROC curve (right) for the GNN interaction type classification method applied to CC and NC neutrino events. The GNN has been trained with charged current interaction as signal, and the GNN model score is the network output. The fiTQun-based model score is calculated as described in 2.2.1.

NC events, as well as a portion of CC events being wrongly identified. But if the only goal is a pure sample of CC events, this is irrelevant.

3.4 Reconstruction speed

Recorded training and reconstruction times for the GNN are shown in table 5 and compared to the average event reconstruction time of the fiTQun framework. The benchmarking was done using the full neutrino data set of 800,000 events.

Table 5. Comparison of reconstruction speeds of the fiTQun-based and GNN classification methods, with and without training of the GNN.

| | |
|-----------------------|--|
| Data extraction | 10^{-4} mins/event |
| Training | 10^{-3} mins/event |
| Reconstruction | 10^{-4} mins/event |
| fiTQun Reconstruction | 1 min/event |
| Improvement | 10^3 (w/ training) / 10^4 (w/o training) |

Using the GraphNeT framework involves some overhead when converting data to one of the database formats, for which the framework is optimised. But the majority of the time consumption lies in the training step. Including model training, the total reconstruction is faster than the fiTQun reconstruction by a factor of 10^3 . This trend would only improve with more events, since one only needs to train on a subset of the events, the size of which can be kept constant.

In a scenario where one could use an already trained model to reconstruct events (for instance if one makes new simulations with only minor changes to the detector layout), the reconstruction speed improvement can reach a factor of 10^4 .

The GNN training and inference is performed on a GPU, with a CPU handling the data loading and graph building. The fiTQun-based reconstruction is run on a CPU. A GPU version of fiTQun has been in development, but was not available at the time of this study [31]. Due to the simplicity of the operations involved in ML inference, running the GNN on a GPU takes advantage of the parallelisation the GPU offers. Thus, a direct speed comparison is hard to perform, but even considering differences in hardware, a speed-up of 10^3 is a significant gain.

4 Conclusion

In this work we propose a new classification method for the ESS ν SB based on GNNs and the GraphNeT framework [20]. We demonstrate that the GNN-based method has comparable performance to the fiTQun method on simulated charged lepton events with likelihood-based cuts applied to reject difficult events, and that the GNN performance surpassed that of fiTQun for simulated CC neutrino events both with and without the aforementioned cuts.

We also present results that indicate that the GNN-based method is able to accurately reconstruct more events with two Cherenkov rings (from muons decaying to electrons), and demonstrate how GNNs can be trained to separate CC and NC events as well as events with and without pion production, and how this is beneficial, as the event types differ in reconstruction performance.

The reconstruction carried out using the GNN method is demonstrated to have a factor of 10^4 increase in performance speed from the fiTQun method. This will be very beneficial to the ESS ν SB experiment, providing great flexibility in the detector design phase.

Acknowledgments

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

The authors acknowledge support provided by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska -Curie grant agreement No 860881-HIDDeN.

The training of models and reconstruction of events were performed using the GraphNeT framework (Apache 2.0), developed and maintained by GraphNeT group. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) and the Swedish National Infrastructure for Computing (SNIC) at LUNARC partially funded by the Swedish Research Council through grant agreements no. 2022-06725 and no. 2018-05973.

The authors acknowledge the assistance from C. Vilela, E. O’Sullivan, H. Tanaka, B. Quilain and M. Wilking for the use of the WCSim and fitQun software packages.

A Events displays of events with superior GNN Performance

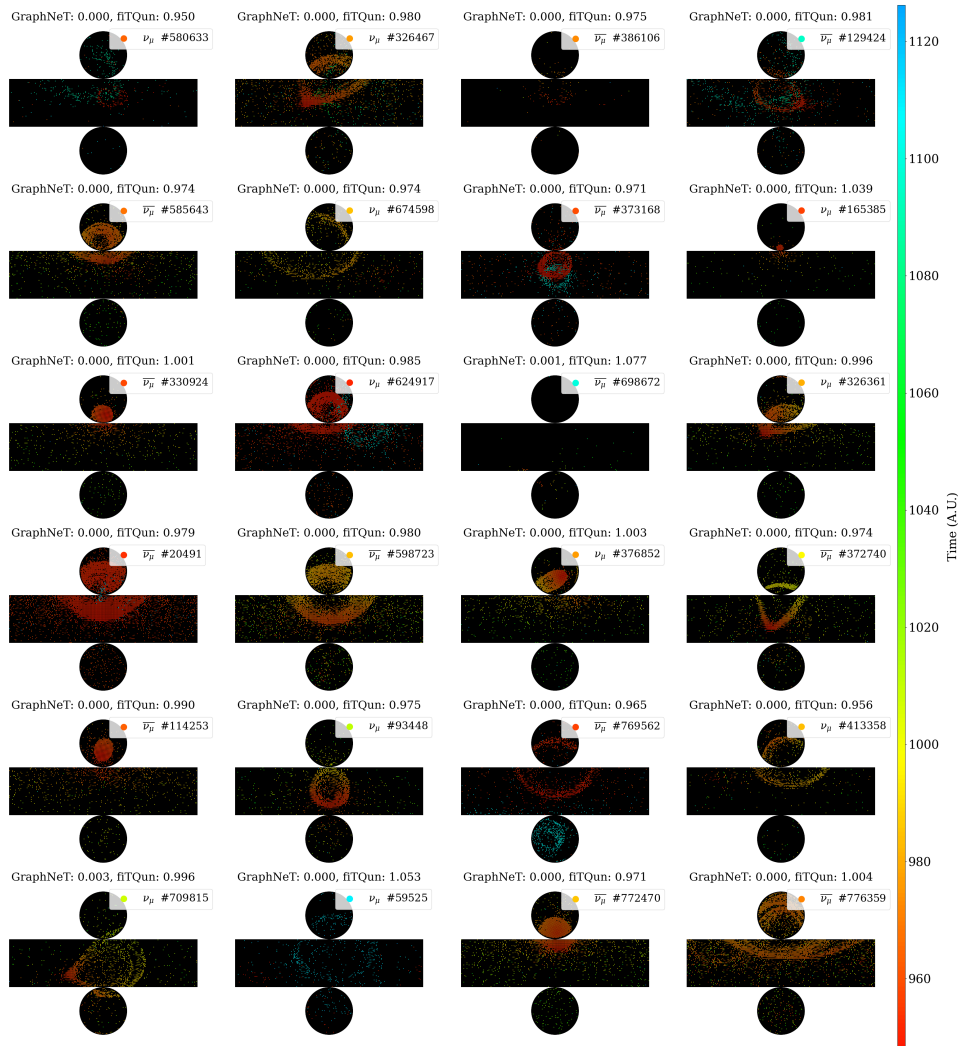


Figure 25. Events displays of randomly selected CC interaction neutrino events for which the GNN flavour classification method predicts correctly with great confidence, while the fitQun-based method predicts incorrectly with great confidence. Several of the event displays show multiple rings of Cherenkov light, indicating that the GNN has the ability to accurately classify low energy muon neutrino events.

References

- [1] A. Cervera et al., *Golden measurements at a neutrino factory*, *Nucl. Phys. B* **579** (2000) 17 [[hep-ph/0002108](#)].
- [2] P. Coloma and E. Fernandez-Martinez, *Optimization of neutrino oscillation facilities for large θ_{13}* , *JHEP* **04** (2012) 089 [[arXiv:1110.4583](#)].
- [3] R. Garoby et al., *The European Spallation Source Design*, *Phys. Scripta* **93** (2018) 014001.
- [4] A. Alekou et al., *The European Spallation Source neutrino super-beam conceptual design report*, *Eur. Phys. J. ST* **231** (2022) 3779 [[arXiv:2206.01208](#)].
- [5] ESSnUSB collaboration, *The ESSnUSB Design Study: Overview and Future Prospects*, *Universe* **9** (2023) 347 [[arXiv:2303.17356](#)].
- [6] T. Fukuda et al., *First neutrino event detection with nuclear emulsion at J-PARC neutrino beamline*, *PTEP* **2017** (2017) 063C02 [[arXiv:1703.03659](#)].
- [7] IceCube collaboration, *Low energy event reconstruction in IceCube DeepCore*, *Eur. Phys. J. C* **82** (2022) 807 [[arXiv:2203.02303](#)].
- [8] T2K collaboration, *Improving the T2K Oscillation Analysis with fitQun: A New Maximum-Likelihood Event Reconstruction for Super-Kamiokande*, *J. Phys. Conf. Ser.* **888** (2017) 012066.
- [9] R. Abbasi et al., *Graph Neural Networks for low-energy event classification & reconstruction in IceCube*, *2022 JINST* **17** P11003 [[arXiv:2209.03042](#)].
- [10] A. Burgman, J. Park, J. Cederkall and P. Christiansen, *The ESS Neutrino Super-Beam Near Detector*, *PoS EPS-HEP2021* (2022) 797 [[arXiv:2111.05550](#)].
- [11] C. Andreopoulos et al., *The GENIE Neutrino Monte Carlo Generator*, *Nucl. Instrum. Meth. A* **614** (2010) 87 [[arXiv:0905.2517](#)].
- [12] C. Andreopoulos et al., *The GENIE Neutrino Monte Carlo Generator: Physics and User Manual*, [arXiv:1510.05494](#).
- [13] E. O’Sullivan, *The WCSim GEANT4 application*, <https://github.com/WCSim/WCSim>.
- [14] GEANT4 collaboration, *GEANT4 — A Simulation Toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250.
- [15] SUPER-KAMIOKANDE collaboration, *Atmospheric Neutrino Oscillation Analysis with Improved Event Reconstruction in Super-Kamiokande IV*, *PTEP* **2019** (2019) 053F01 [[arXiv:1901.03230](#)].
- [16] R.B. Patterson et al., *The Extended-track reconstruction for MiniBooNE*, *Nucl. Instrum. Meth. A* **608** (2009) 206 [[arXiv:0902.2222](#)].
- [17] R. Abbasi et al., *A Convolutional Neural Network based Cascade Reconstruction for the IceCube Neutrino Observatory*, *2021 JINST* **16** P07041 [[arXiv:2101.11589](#)].
- [18] T. Zhu, M. Jin and C.A. Argüelles, *Comparison of geometrical layouts for next-generation large-volume Cherenkov neutrino telescopes*, *2025 JINST* **20** P05001 [[arXiv:2407.19010](#)].
- [19] F. Scarselli et al., *The Graph Neural Network Model*, *IEEE Trans. Neural Networks* **20** (2009) 61.
- [20] A. Søgaard et al., *GraphNet: Graph neural networks for neutrino telescope event reconstruction*, *J. Open Source Softw.* **8** (2023) 4971 [[arXiv:2210.12194](#)].
- [21] IceCube collaboration, *The Design and Performance of IceCube DeepCore*, *Astropart. Phys.* **35** (2012) 615 [[arXiv:1109.6096](#)].
- [22] IceCube collaboration, *The IceCube Neutrino Observatory: Instrumentation and Online Systems*, *2017 JINST* **12** P03012 [Erratum *ibid.* **19** (2024) E05001] [[arXiv:1612.05093](#)].

- [23] A. Paszke et al., *PyTorch: An Imperative Style, High-Performance Deep Learning Library*, in the proceedings of the *33rd International Conference on Neural Information Processing Systems*, Vancouver, Canada, December 8–14 (2019) [[arXiv:1912.01703](#)].
- [24] M. Fey and J.E. Lenssen, *Fast Graph Representation Learning with PyTorch Geometric*, in the proceedings of the *ICLR Workshop on Representation Learning on Graphs and Manifolds*, New Orleans, LA, USA, May 6 (2019) [[arXiv:1903.02428](#)].
- [25] W. Falcon et al., *Pytorchlightning/pytorch-lightning: 0.7.6 release*, DOI:10.5281/zenodo.3828935.
- [26] Y. Wang et al., *Dynamic Graph CNN for Learning on Point Clouds*, *ACM Trans. Graphics* **38** (2019) 1.
- [27] L.N. Smith, *A disciplined approach to neural network hyper-parameters: Part 1 — learning rate, batch size, momentum, and weight decay*, Tech. Rep. 5510-026, US Naval Research Laboratory (2018) [[arXiv:1803.09820](#)].
- [28] L.N. Smith and N. Topin, *Super-convergence: very fast training of neural networks using large learning rates*, *Proc. SPIE* **11006** (2019) 1100612.
- [29] PARTICLE DATA GROUP collaboration, *Review of particle physics*, *Phys. Rev. D* **110** (2024) 030001.
- [30] D. Rein and L.M. Sehgal, *Neutrino Excitation of Baryon Resonances and Single Pion Production*, *Annals Phys.* **133** (1981) 79.
- [31] R. Calland, *Accelerating fitqun with gpus*, presented at *fitQun Workshop at TRIUMF*, Vancouver, BC, Canada, August 18–20 (2015).

Author List

J. Aguilar^a, M. Anastasopoulos^b, D. Barčot^r, E. Baussan^c, A.K. Bhattacharyya^b, A. Bignami^b, M. Blennow^{d,e}, M. Bogomilov^f, B. Bolling^b, E. Bouquerel^c, F. Bramati^g, A. Branca^g, G. Brunetti^g, A. Burgman^{h,1}, I. Bustinduy^a, C.J. Carlile^h, J. Cederkall^h, T.W. Choiⁱ, S. Choubey^{d,e}, P. Christiansen^h, M. Collins^{j,b}, E. Cristaldo Morales^g, P. Cupia^k, D. D'Ago^y, H. Danared^b, J.P.A.M. de André^c, M. Dracos^c, I. Efthymiopoulos^l, T. Ekelöfⁱ, M. Eshraqi^b, G. Fanourakis^m, A. Farrickerⁿ, E. Fasoula^{o,p}, T. Fukuda^q, J. García-Marcos^u, N. Gazis^b, Th. Geralis^m, M. Ghosh^r, A. Giarnetti^s, G. Gokbulut^{t,u}, C. Hagner^v, L. Halić^r, M. Hooft^u, K.E. Iversen^{h,*}, N. Jachowicz^u, M. Jenssen^b, R. Johansson^b, I. Karakoulias^{m,o}, E. Kasimi^o, A. Kayis Topaksu^t, B. Kildetoft^b, B. Kliček^r, K. Kordas^{o,p}, A. Leisos^w, M. Lindroos^{b,h,2}, A. Longhin^x, C. Maiano^b, S. Marangoni^g, C. Marrelli^l, D. Meloni^s, M. Mezzetto^y, N. Milas^b, J.L. Muñoz^a, K. Niewczas^u, M. Oglakci^t, T. Ohlsson^{d,e}, M. Olvegårdⁱ, M. Pari^w, J. Park^{h,3}, D. Patrzałek^b, G. Petkov^f, Ch. Petridou^{o,p}, P. Poussot^c, A. Psallidas^m, F. Pupilli^y, D. Saiang^z, D. Sampsonidis^{o,p}, A. Scanu^g, C. Schwab^c, F. Sordo^a, A. Sosa^{aa}, G. Stavropoulos^m, M. Stipčević^r, R. Tarkeshian^b, F. Terranova^g, T. Tolba^v, E. Trachanas^b, R. Tsenov^f, A. Tsirigotis^w, S.E. Tzamaras^o, M. Vanderpoorten^u, G. Vankova-Kirilova^f, N. Vassilopoulos^{ab}, S. Vihonen^{d,e}, J. Wurtz^c, V. Zeter^c, O. Zormpa^m

^a Consorcio ESS-bilbao, Parque Científico y Tecnológico de Bizkaia, Laida Bidea, Edificio 207-B, 48160 Derio, Bizkaia, Spain

^b European Spallation Source, Box 176, SE-221 00 Lund, Sweden

^c IPHC, Université de Strasbourg, CNRS/IN2P3, Strasbourg, France

^d Department of Physics, School of Engineering Sciences, KTH Royal Institute of Technology, Roslagstullsbacken 21, 106 91 Stockholm, Sweden

^e The Oskar Klein Centre, AlbaNova University Center, Roslagstullsbacken 21, 106 91 Stockholm, Sweden

^f Sofia University St. Kliment Ohridski, Faculty of Physics, 1164 Sofia, Bulgaria

^g University of Milano-Bicocca and INFN Sez. di Milano-Bicocca, 20126 Milano, Italy

^h Department of Physics, Lund University, P.O. Box 118, 221 00 Lund, Sweden

ⁱ Department of Physics and Astronomy, FREIA Division, Uppsala University, P.O. Box 516, 751 20 Uppsala, Sweden

^j Faculty of Engineering, Lund University, P.O. Box 118, 221 00 Lund, Sweden

^k AGH University of Krakow, al. A. Mickiewicza 30, 30-059 Krakow, Poland

^l CERN, 1211 Geneva 23, Switzerland

^m Institute of Nuclear and Particle Physics, NCSR Demokritos, Neapoleos 27, 15341 Agia Paraskevi, Greece

ⁿ Cockcroft Institute (A36), Liverpool University, Warrington WA4 4AD, U.K.

^o Department of Physics, Aristotle University of Thessaloniki, Thessaloniki, Greece

^p Center for Interdisciplinary Research and Innovation (CIRI-AUTH), Thessaloniki, Greece

^q Institute for Advanced Research, Nagoya University, Nagoya 464–8601, Japan

^r Center of Excellence for Advanced Materials and Sensing Devices, Ruđer Bošković Institute, 10000 Zagreb, Croatia

^s Dipartimento di Matematica e Fisica, Università di Roma Tre, Via della Vasca Navale 84, 00146 Rome, Italy

^t University of Cukurova, Faculty of Science and Letters, Department of Physics, 01330 Adana, Turkey

^u Department of Physics and Astronomy, Ghent University, Proeftuinstraat 86, B-9000 Ghent, Belgium

^v Institute for Experimental Physics, Hamburg University, 22761 Hamburg, Germany

^w Physics Laboratory, School of Science and Technology, Hellenic Open University, 26335, Patras, Greece

^x Department of Physics and Astronomy “G. Galilei”, University of Padova and INFN Sezione di Padova, Italy

^y INFN Sez. di Padova, Padova, Italy

^z Department of Civil, Environmental and Natural Resources Engineering Luleå University of Technology, SE-971 87 Lulea, Sweden

^{aa} ISIS, STFC, Rutherford Appleton Laboratory, Harwell Oxford, Didcot OX11 0QX, United Kingdom

^{ab} Institute of High Energy Physics (IHEP) Dongguan Campus, Chinese Academy of Sciences (CAS), Guangdong 523803, China

* Corresponding author

¹ *Current address: Department of Physics, Stockholm University, 106 91 Stockholm, Sweden.*

² *Deceased.*

³ *Current address: Center for Exotic Nuclear Studies, Institute for Basic Science, 34126 Daejeon, Republic of Korea.*