

Neutrino characterisation using convolutional neural networks in CHIPS water Cherenkov detectors

Josh Tingey,^{a,*} Simeon Bash,^a John Cesar,^a Thomas Dodwell,^a Stefano Germani,^a Paul Kooijman,^{b,c} Petr Mánek,^a Mustafa Ozkaynak,^a Andy Perch,^a Jennifer Thomas^{a,d} and Leigh Whitehead^a

^aDepartment of Physics and Astronomy, UCL,
Gower St, London, WC1E 6BT, U.K.

^bNikhef National Institute for Subatomic Physics,
Science Park 105, 1098 XG, Amsterdam, The Netherlands

^cUniversity of Utrecht,
Princetonplein 5 3584 CC Utrecht, The Netherlands

^dDepartment of Physics, University of Wisconsin Madison,
1150 University Ave, Madison, Wisconsin, U.S.A.

E-mail: j.tingey.16@ucl.ac.uk

ABSTRACT: This work presents a novel approach to water Cherenkov neutrino detector event reconstruction and classification. Three forms of a Convolutional Neural Network have been trained to reject cosmic muon events, classify beam events, and estimate neutrino energies, using only a slightly modified version of the raw detector event as input. When evaluated on a realistic selection of simulated CHIPS-5kton prototype detector events, this new approach significantly increases performance over the standard likelihood-based reconstruction and simple neural network classification.

KEYWORDS: Cherenkov detectors; Neutrino detectors; Particle identification methods

ARXIV EPRINT: [2206.14904](https://arxiv.org/abs/2206.14904)

*Corresponding author.

Contents

1	Introduction	1
1.1	The CHIPS R&D project	2
1.2	Modern machine learning for CHIPS	3
2	Detector simulation	4
3	Convolutional neural networks for CHIPS	5
3.1	Inputs	6
3.2	Architecture	7
3.3	Outputs	9
3.3.1	Outputs — cosmic rejection	10
3.3.2	Outputs — beam classification	10
3.3.3	Outputs — energy estimation	11
3.4	Training	11
3.4.1	Training — cosmic rejection	12
3.4.2	Training — beam classification	13
3.4.3	Training — energy estimation	13
4	Results	14
4.1	Evaluation sample	14
4.2	Preselection cuts	14
4.3	Cosmic rejection and containment	15
4.4	CC ν_e selection	18
4.5	CC ν_μ selection	20
4.6	Interaction type classification	23
4.7	Energy estimation	23
4.8	Comment	26
5	Conclusions and future work	27

1 Introduction

In pursuit of answers to some of the open questions in physics, the study of the phenomenon of neutrino oscillations has become paramount, owing to the implication of the neutrinos' small masses in relation to the mass of the Universe and, potentially, the question of the matter anti-matter asymmetry. The next generation of giant neutrino detectors aim to push the size and therefore detector mass to the 100-500 kilo-ton level. Such massive detectors will quickly become limited by systematic uncertainties, whereas most present experiments are still statistics limited.

For detectors to remain practical and affordable into the future, a novel design strategy could be utilized in order to reduce the systematic limitations of single large detectors. Typically systematic limitations due to cross-section uncertainties, neutrino event modelling and even flux uncertainties from the neutrino beam can be alleviated by having functionally identical near and far detectors such that these systematics cancel between them.

However, if the detector cost, which is presently on the order of \$20M/kilo-ton, can be reduced such that several identical far detectors can be afforded, and then placed at different positions and distances in the neutrino beam, the oscillation parameters could be measured with exact cancellation of the aforementioned systematic errors. In addition, systematic errors which do not cancel in the typical two-detector setup, for example, the electron neutrino cross-section, which is important in only the far detector in the traditional two-detector setup will cancel in the multiple far detector setup.

1.1 The CHIPS R&D project

The CHIPS R&D project aims to develop such a novel strategy: a ‘cheap as chips’ water Cherenkov detector has been designed [1] and was built in 2018-2019. Primarily aimed for deployment in long-baseline accelerator beam scenarios, the idea was to lower the cost per kt of sensitive mass to between \$200k-\$300k. For comparison, the Super-Kamiokande detector cost approximately \$4 million per kt to build more than 30 years ago. As physics sensitivity depends on the detector performance in addition to sensitive mass, this comparison is not entirely rigorous; however, it highlights the scale of possible cost savings. This paper focuses on the simulation and reconstruction of events in the 5 kton CHIPS detector (CHIPS-5) 7 milli-radians off-axis in the NuMI beam [2] at a baseline of 710 kilometers, in northern Minnesota.

There are several components to the cost savings: the water acts as a mechanical support and cosmic overburden, thereby saving costs of excavation and structural material. The electronics are developed using off-the-shelf technology where possible, saving significantly not only on engineering but also on cables and HV infrastructure. Finally, the layout and choice of the PMTs saves on photocathode area with very little loss in efficiency for neutrino events in the 1-10 GeV range.

The CHIPS detector module is a cylindrical water Cherenkov detector submerged in a deep body of water on the Earth’s surface, such as a lake, reservoir, or as in the first prototype, a flooded mine pit. Figure 1 shows a graphical rendering of a typical CHIPS detector module. The instrumentation on the walls and endcaps of the detector geometry is visible, a detailed description of which is given later.

However, such bodies of water are not available everywhere, and some compromise between baseline distance and off-axis angle must be incorporated to maximize the sensitivity.

The common practice of using majority bespoke components is replaced by using commercially available components wherever possible. The number of expensive elements, such as photomultiplier tubes, are also reduced by only considering multi-GeV accelerator beam neutrino events, such that full high-density detector instrumentation is not required.

The CHIPS detectors are not only designed to be cheap but also practical. Straightforward to build, quick to deploy, and upgradable once operational, multiple detector modules could be flexibly combined depending on available resources and funding. Compared to DUNE and Hyper-Kamiokande, both of which require a large upfront budget and many years to construct, cheap CHIPS detector modules could be deployed as needed in under two years by a relatively small team.




Figure 1. Graphical rendering of the fully deployed and submerged CHIPS-5 detector module with a section of the liner cutaway. The bottom endcap and wall instrumentation is visible, as well as the top endcap structure and floatation. The faint green lines indicate the cables holding the endcaps together.

1.2 Modern machine learning for CHIPS

Alongside work to realise the concept outlined above, much effort has been expended to explore alternative water Cherenkov neutrino event reconstruction and classification techniques for CHIPS style detectors. This work is motivated by the need to maximise the achievable physics performance per unit cost.

Over the last few years, neutrino experiments have adopted modern machine learning techniques for a range of event analysis tasks [3]. In 2016 the NOvA experiment applied a Convolutional Neural Network (CNN) [4], a type of deep learning [5] neural network, to the task of classifying the interaction type of events within their sampling calorimeter detector [6]. Two views of raw detector data were used as input to train a CNN network [7]. Further NOvA iterations have since been applied to both the classification of individual energy deposit clusters [8] and ν_e and e^- energy reconstruction [9].

The IceCube collaboration, who also make use of Cherenkov radiation [10], utilize deep neural networks to distinguish between cascade events (CC ν_e) and muon tracks (CC ν_μ) as they traverse the ice.

CNNs have been applied to water Cherenkov detector events by both the Daya Bay [11] and KM3NeT/ORCA collaborations [12].

CNNs have also been applied to liquid argon time-projection chambers. The MicroBooNE experiment [13] has shown that in addition to classification tasks and particle identification [14], the spatial localisation of single particles within events is also possible [15]. The DUNE collaboration has designed a network to output both the interaction class and counts of different particle types

within an event [16, 17], as well as kinematic energy estimation [18]. This approach is called *multi-task* learning and is discussed in detail in this paper in section 3.3. Raw data de-noising, the first step of reconstruction, is explored in ProtoDUNE [19], with Graph Neural Networks also being investigated as an alternative to CNNs.

Sub-manifold sparse convolutional networks (SSCNs) [20], used when there is sparse input data, have been suggested to account for the poor performance of naive applications of CNNs in large scale liquid argon time-projection chambers such as DUNE. The SSCNs can be used for image classification and object detection in such detectors and can have over 95% pixel clustering efficiency and purity for Michel electrons [21]. There are also investigations of the efficiencies of different CNN based algorithms such as Residual Networks and Inception Networks algorithms reconstructing shower energies in a LArTPC [22].

The output from the CHIPS detector is a simple *image* of each event where two pieces of information are known for each photomultiplier tube (PMT): the number of collected photo-electrons and the associated hit times. CNNs primarily developed for image-based computer vision tasks are analogous.

This work outlines an event analysis methodology for CHIPS following this natural progression. Three forms of CNN have been developed: one for cosmic muon rejection, one for beam event classification and one for neutrino energy estimation, all using only a slightly modified version of the raw detector event as input. The detector simulation is first described and a description of how CNNs are used for CHIPS is presented. The new methodology's performance is then detailed. For evaluation purposes, only the implementation as applied to the CHIPS-5 prototype detector module is considered in this work.

2 Detector simulation

The detector simulation used throughout this work [23] builds upon the WCSim water Cherenkov simulation package [24] which employs the Geant4 simulation framework [25–27]. Developed initially to simulate a large water Cherenkov detector in the Long Baseline Neutrino Experiment beam [28], WCSim is now used more widely in the field.

The simulation builds an n -sided, regular polygonal prism consisting of two endcaps and a barrel, filled with water and lined with a low reflectivity *blacksheet*. The geometry is separated into *regions* within both the barrel and endcaps. Each region is filled with a unique base unit of geometry known as the *unit cell*.

The unit cell defines a pattern of any number of PMTs, including their relative positions and in which direction they face. The final geometry is built by tiling the defined regions with their respective unit cell scaled to match the required regional photocathode coverage. Although exact detector PMT positions are not explicitly defined using this procedure, a given configuration will deterministically generate the same geometry.

In this work, the CHIPS-5 geometry is generated with 28 sides and regions matching the photocathode coverage of the CHIPS-5 detector design. A conservative photon attenuation length of 50 m at 405 nm is used alongside negligible scattering [29], with the PMT glass reflectivity set to 24% [30], and the blacksheets reflectivity kept at the WCSim default of 4%. Note that veto PMTs are ignored for simplicity.

Using the expected NuMI flux [31] at CHIPS-5, the GENIE neutrino event generator (version 3.0.6) [32, 33] is used to generate beam neutrino events using default neutrino cross-sections on water. All channels of neutrino interaction are considered and can be broken down into five main categories:

- **Quasi-Elastic scattering (QEL):** the dominant channel for energies below 1 GeV, involving the neutrino scattering off the entire nucleon.
- **Meson Exchange Current (MEC):** additional contribution channel for energies below 1 GeV, involving two nucleons and producing two protons in the final state.
- **Resonant pion production (Res):** the dominant channel between 1 and 2 GeV, involving the neutrino exciting the nucleon into a resonant state.
- **Coherent pion production (Coh):** a mechanism where the neutrino scatters coherently from the entire nucleus.
- **Deep Inelastic Scattering (DIS):** the dominant channel for neutrino energies above 3 GeV with the additional energy allowing for the neutrino to resolve the individual quark content of the nucleon.

The Cosmic-Ray Shower Library (CRY) [34, 35] is used for cosmic ray event generation, assuming a CHIPS-5 overburden of 50 m and a 2.2 MeV/cm^2 muon energy loss in water as suggested by [36]. Beam event vertices, defined as the insertion position in the simulation, are uniformly randomly placed within the inner detector volume. Cosmic event vertices are inserted into the simulation volume 1 m above the detector volume.

WCSim then simulates the passage of all particles through the detector materials, with interactions, decays, and Cherenkov emission all considered. Whenever a photon is calculated to have hit the photocathode of a PMT, an angular dependent acceptance efficiency is applied to see if it is recorded [30]. If accepted, all hits within 200 ns windows are grouped to form a single recorded hit, with a smeared first hit time used as the recorded time. The standard WCSim methodology is used to determine the total output charge of the hit, given the number of incident photons. This procedure involves a single photoelectron charge distribution being repeatedly probed for each photon before the combined sum is returned which corresponds to the sum of the simulated responses for each photon [37]. The simulation output, used as the input for the rest of this work, is a collection of hits for each neutrino interaction event, each with an associated number of photoelectrons and hit time.

3 Convolutional neural networks for CHIPS

For the majority of HEP experiments, event analysis entails the separation of signal from background, the identification of particle types, the determination of spatial properties, and the estimation of energies. The same is true for the CHIPS detectors, with the primary aims being the selection of CC ν_e signal beam events from a sizeable background (beam and cosmic) and the estimation of the associated neutrino energies.

For this purpose, the CHIPS reconstruction originally relied upon a likelihood-based reconstruction algorithm and a simple classification neural network driven by hand-engineered features [38].

Both suffer from only considering what has been implemented in software and consequently what features are explicitly extracted and then modelled from the data. This restriction makes them prone to ignoring the wide range of edge cases not contained within the bulk of neutrino events and unable to use all the underlying informative features of the data.

In the new methodology, the *image*-like nature of the raw detector output is a natural fit for the standard grid-like input to a CNN. Notably, using a CNN on the raw detector event removes the requirement to build hand-engineered features as the CNN learns to extract the most powerful features within the input data itself [39]. Neutrino events where it is hard to determine the typically reconstructed features, such as neutral current events or events where multiple particles are involved, benefit from the generalisability of the CNN. Three forms of a baseline CNN have been developed to achieve the primary aims outlined above:

- The **cosmic rejection** form aims to prevent the vast cosmic muon background from contaminating the final selected sample of beam events. Therefore, the primary task is a simple binary classification between beam-like and cosmic-like events.
- The **beam classification** form aims to separate beam events by their neutrino and interaction type to primarily select a pure and efficient sample of appearing CC ν_e events, but also a sample of survived CC ν_μ events. Therefore, the principal task is a categorical classification between CC ν_e , CC ν_μ , and background Neutral Current (NC) events.
- The **neutrino energy estimation** form aims to accurately estimate the energy of the ν_e and ν_μ signal events. Therefore, the primary task is a regression on the interactive neutrino energy.

A Python-based software package named *chipsnet* [40] has been built for this purpose. By using the high-level *Application Programming Interfaces* (APIs) provided by the Tensorflow framework (version 2.3.0) [41], a complete pipeline including data preparation, network training, and performance evaluation has been implemented. We separate our methodology into four sections: the inputs used, the CNN architecture, the CNN outputs, and finally the training procedure. For each, we detail the differences between the three network forms.

3.1 Inputs

The primary difficulty in applying CNNs to CHIPS (or any cylindrical detector) is determining how to map an event captured on a cylindrical surface to a two-dimensional grid. This must be done in such a way that the underlying Cherenkov emission topology is not distorted, which could inhibit network learning. This work builds upon the ideas outlined in reference [42]. An event is mapped onto a two-dimensional grid as though it is viewed from its estimated interaction vertex position in order to remove any detector shape effects and focus on the underlying event topology and Cherenkov emission profiles.

To estimate the interaction vertex position, the PMT hits of an event are first sliced in both space and time. Gaps in the time ordering of hits are used to separate the event into time slices. Each of these slices then undergoes basic clustering to remove outlying hits and ensure only the dominant collections of hits are considered. Each cleaned slice is then run through a simple geometric vertex finding algorithm to estimate the interaction vertex position. A circular Hough




Figure 2. Three event map representations of a CC resonant ν_e event expected within CHIPS-5. Initiated by a ν_e of energy 3.3 GeV the final state particles above the Cherenkov threshold include an e^- of energy 2.8 GeV and a 0.3 GeV π^0 .

transform algorithm, traditionally used for water Cherenkov rings is then applied [43]. As output, the voting-based transformation produces a space within which rings of PMT hits exist as peaks. The interaction vertex position is further refined using the Hough peaks in this space.

Using θ and ϕ components calculated as viewed from the estimated interaction vertex position facing along the beam axis, hit PMTs are mapped onto a 64×64 grid. This procedure is used to generate two event *maps*. A *hit-charge* map where each grid bin is given by the total collected photoelectrons from all PMTs mapped to that bin and a *hit-time* map where each grid bin is given by the first hit time (in nanoseconds) across all PMTs mapped to that bin. Each hit-time map is further corrected so that the first hit time across all bins lies at zero.

By design, the Hough transform uses the estimated interaction vertex position to generate the transform space. Therefore, by re-binning the transform space to a 64×64 grid, a third *Hough-height* map is generated for each event. This event map provides a complementary but different representation of the event where Cherenkov rings are instead represented as peaks, allowing for additional discriminating features to be learnt.

All three event maps: hit-charge, hit-time, and Hough-height, are down-sampled using 8-bit encoding by converting each 32-bit float value to an integer between 0 and 255. Encoding significantly reduces storage requirements and dramatically increases the speed with which data can be loaded during training (which is the primary training bottleneck). For each map type, a range over which to encode from zero up to a *cap-point* is chosen to minimise the number of bin values that are capped at the maximum encoded value of 255. The generated event maps for an example CC ν_e event are shown in figure 2. Each network form (cosmic rejection, beam classification, and energy estimation) uses the same event maps as input.

3.2 Architecture

An illustrative diagram of the baseline chipsnet CNN architecture is shown in figure 3. Based on the Visual Geometry Group (VGG) network [44] there are a few key differences from the literature defined network:

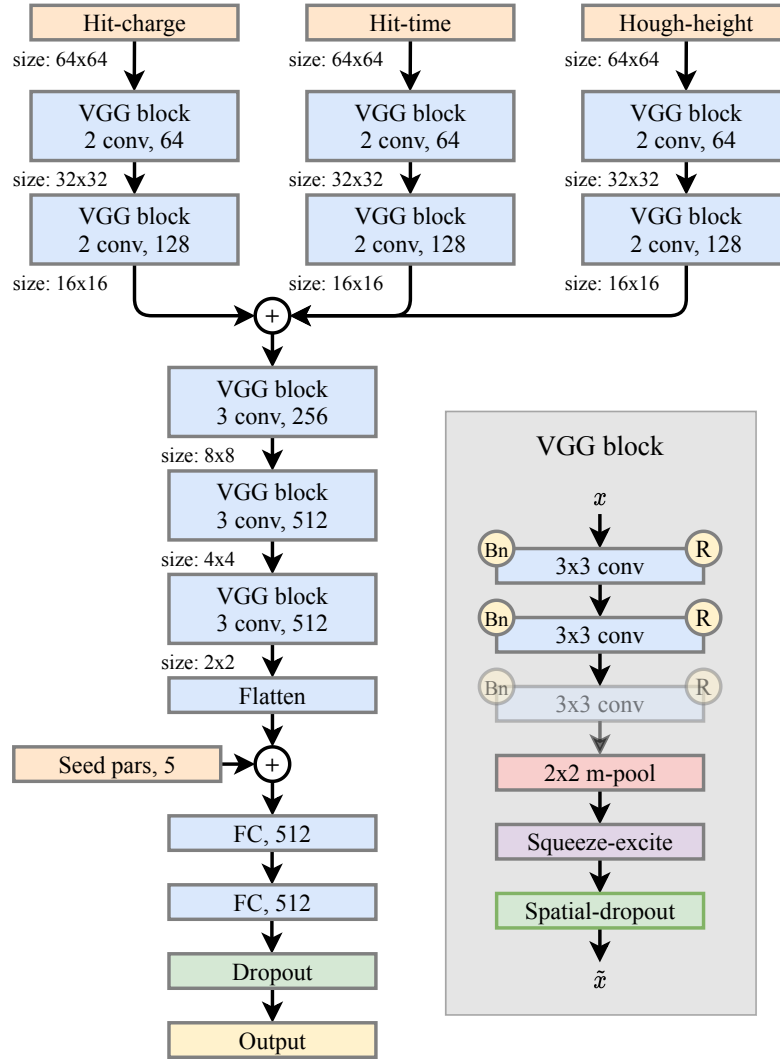


Figure 3. Illustrative diagram of the baseline chipsnet architecture. The three input event maps are separately passed through two VGG blocks each before their outputs are combined and passed through a further three VGG blocks together. The flattened VGG blocks outputs are then concatenated with five seed parameters (seed pars) and passed through two fully-connected (FC) layers of 512 neurons each before the output layer. Both the number of convolutional units (1st value) and kernels (2nd value) is shown for each block. The detailed VGG block structure is shown within the grey box. The circular yellow R and Bn indicate the use of the ReLU activation function and batch normalisation, respectively.

- Each of the three event maps: hit-charge, hit-time, and hough-height, are initially fed into three separate branches. Each branch contains two VGG blocks with two convolutional (conv) layers each (four conv layers in total). The outputs from each branch are merged using a concatenation layer before being fed to the rest of the network. This configuration allows for event map specific features to be learnt independently before combined features are learnt by the rest of the network.
- Batch normalisation [45] is included before the activation (ReLU) function for every conv layer.

- Squeeze-and-excitation units [46] are included after the max-pooling operation in all VGG blocks. These units introduce extra parameters to model the interdependencies between output feature maps, allowing the network to learn how to weight each feature map effectively.
- Dropout is included at the end of each VGG block as well as after the final fully-connected layer. Instead of dropping individual kernel elements, the dropout within the VGG blocks drops entire kernels at each training iteration, this is commonly called *two-dimensional spatial dropout*. The dropout after the fully-connected layers is standard, in that it drops out individual fully-connected neurons.
- Five parameters from the estimation of the interaction vertex position (seed pars) are concatenated with the flattened layer before the fully-connected layers. Included are the three components of the estimated interaction vertex position (s_x, s_y, s_z) , and the two components of the estimated track direction (s_θ, s_ϕ) . These parameters provide the network with spatial context as to where the input event maps have been generated in the detector and the dominant direction of PMT activity.

The chipsnet baseline architecture is implemented using the Keras API built into TensorFlow [47]. Each network form (cosmic rejection, beam classification, and energy estimation) uses the same baseline model architecture.

3.3 Outputs

Many CNN applications are found to benefit from learning multiple tasks simultaneously. This seems to be because training with multiple tasks tends to return a network with an improved generalised representation of the inputs, with features learnt for one task improving the performance of another. Increasing the number of tasks also acts as a regularization technique making it harder for a CNN to overfit on any one single task. Commonly referred to as *multi-task* learning, this methodology is used here.

To train a network with multiple tasks (outputs), a loss function E_{tot} , must be defined to combine the individual loss functions for each task E_i . The simplest way to do this is via a linear weighted sum, such that

$$E_{\text{tot}} = \sum_{i=1}^{i=N} w_i E_i, \quad (3.1)$$

where N is the number of tasks and w_i are the associated weights. In this work this is referred to as the *simple* multi-task loss.

The final network performance can strongly depend on the relative weighting between loss functions, especially when the values returned by each differ by many orders of magnitude (common when combining regression and classification tasks). Therefore, finding the optimal w_i weights can be both difficult and time-consuming. Another approach outlined in reference [48] remedies this problem by learning the optimal weighting between loss functions. This is done by introducing an additional trainable parameter σ_i , for each loss function, such that

$$E_{\text{tot}} = \sum_{i=1}^{i=N} \frac{1}{2\sigma_i^2} E_i + \log \sigma_i. \quad (3.2)$$

In this work we refer to this as the *learnt* multi-task loss.

The specific number and nature of outputs for the specific network forms are detailed below. Although physically motivated to some degree, the exact set of tasks and the loss combination technique used is mainly driven by extensive trial-and-error.

3.3.1 Outputs — cosmic rejection

Alongside the primary task of classification between beam-like and cosmic-like events, training the network to also separate events where the primary charged lepton does or does not escape the detector volume is found to improve cosmic rejection performance. As a large proportion of cosmic muons are relatively high in energy and so escape the detector, there is motivation as to why this additional task is helpful. Hence the two outputs for this network form are:

1. **Cosmic score (1 classification neuron):** returns a score between zero and one corresponding to whether the event is beam or cosmic like.
2. **Escapes score (1 classification neuron):** returns a score between zero and one corresponding to whether the charged lepton in an event is contained or escapes the detector. NC beam events without a charged lepton are masked (do not contribute to the loss) during training for this output.

Both outputs are trained using a binary cross-entropy loss function and combined using the *simple* multi-task loss in equation (3.1), each with an arbitrarily chosen weight of 1.

3.3.2 Outputs — beam classification

Alongside the primary task of classification between CC ν_e , CC ν_μ , and background NC events, training the network on additional classification and particle counting tasks is found to improve performance. Note that the particle counting tasks are not used in this work for anything but to increase the primary classification performance. However, future work could exploit any ability to separate exclusive final states, deduced from these particle counts, to improve energy resolution and systematic errors. Hence the nine outputs for this network form are:

1. **Combined category (3 classification neurons):** returns a classification probability score between zero and one for each of CC ν_e , CC ν_μ , and NC (summing to one).
2. **CC category (6 classification neurons):** returns a classification probability score between zero and one for each of CC, Quasi-Elastic (QEL), Deep Inelastic Scattering (DIS), Resonant (Res), Coherent (Coh), Meson Exchange Current (MEC), and CC-other (summing to one). NC events are masked (do not contribute to the loss) during training for this output.
3. **NC category (4 classification neurons):** returns a classification probability score between zero and one for each of NC-Res, NC-DIS, NC-Coh, and NC-other (summing to one). CC events are masked (do not contribute to the loss) during training for this output.
4. **Electron count (4 classification neurons):** returns a classification probability score between zero and one for each of 0, 1, 2, and 3+ electrons in the final state (summing to one).
5. **Muon count (4 classification neurons):** returns a classification probability score between zero and one for each of 0, 1, 2, and 3+ muons in the final state (summing to one).

6. **Proton count (4 classification neurons):** returns a classification probability score between zero and one for each of 0, 1, 2, and 3+ protons in the final state (summing to one).
7. **π^\pm count (4 classification neurons):** returns a classification probability score between zero and one for each of 0, 1, 2, and 3+ charged pions in the final state (summing to one).
8. **π^0 count (4 classification neurons):** returns a classification probability score between zero and one for each of 0, 1, 2, and 3+ neutral pions in the final state (summing to one).
9. **Photon count (4 classification neurons):** returns a classification probability score between zero and one for each of 0, 1, 2, and 3+ photons in the final state (summing to one).

All outputs are trained using a categorical cross-entropy loss function and combined using the *simple* multi-task loss in equation (3.1), each with an arbitrarily chosen weight of 1.

3.3.3 Outputs — energy estimation

Alongside the primary task of estimating the neutrino energy, training the network to additionally estimate the primary charged lepton energy and the interaction vertex position and time are found to improve performance. Although this improvement is relatively small for neutrino energy estimation, it dramatically improves primary charged lepton energy estimation. With two energy tasks, the network is encouraged to learn how the primary charged lepton and neutrino energies are related. As the interaction vertex position within the detector and hence distance from the wall can impact the number of deposited photoelectrons, there is motivation as to why this additional task is also helpful. Hence the six outputs for this network form are:

1. **Neutrino energy (1 regression neuron):** returns the estimated neutrino energy.
2. **Charged lepton energy (1 regression neuron):** returns the estimated charged lepton energy.
3. **Interaction vertex x-position (1 regression neuron):** returns the estimated interaction vertex x-position.
4. **Interaction vertex y-position (1 regression neuron):** returns the estimated interaction vertex y-position.
5. **Interaction vertex z-position (1 regression neuron):** returns the estimated interaction vertex z-position.
6. **Interaction time (1 regression neuron):** returns the estimated interaction time.

All outputs are trained using a mean-squared error loss function and combined using the *learn* multi-task loss in equation (3.2) which includes the additional trainable parameters.

3.4 Training

All networks are trained on an 18 core CPU (36 thread) machine equipped with four NVIDIA GeForce RTX 2080 graphics processing units (GPUs). The Tensorflow dataset API is used to create an efficient input data pipeline where data is loaded on-the-fly at training time. This procedure

ensures all CPU threads are utilised loading, decoding, and preprocessing data for the primary GPU based network calculations before being needed, maximising computational efficiency.

During preprocessing, all 8-bit input event map values are converted to 32-bit float values bounded between zero and one. A random factor scaling is applied to each map bin, generated from a normal distribution centred on one with a standard deviation of σ_r . By fluctuating the bin values the network is forced to focus less on the absolute bin values and more on the underlying event topology. This process provides valuable regularisation to reduce overfitting and makes the trained networks robust to small changes within the input.

A minibatch training strategy of minibatch size, n_b , using the Adam optimiser [49] ($\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10e - 7$) is used. The exact training sample size and composition for each specific network are given below, but for all network forms a 95% training to 5% validation data split is employed across the full training sample. The learning rate for each epoch η_e , is set to decrease throughout training according to

$$\eta_e = \frac{\eta_0}{1 + c_d(e - 1)}, \quad (3.3)$$

where η_0 is the initial learning rate, e is the epoch number (starting at one), and c_d is the learning rate decay coefficient.

When training each network, there is a list of tunable hyperparameters, all of which are optimised using the SHERPA hyperparameter tuning framework [50]. To maximise performance, SHERPA uses a random search algorithm to select random configurations of hyperparameters which are then tested by training the network for five epochs on the available training data. Each configuration's performance is assessed by using a metric detailed for each network below. The search space is confined to a specific range or selection of choices for each hyperparameter, with:

- the **initial learning rate** η_0 , in a range from 0.00005 to 0.001;
- the **learning rate decay coefficient** c_d , in a range from 0.2 to 0.8;
- the **dropout probability** p_d , in a range from 0.0 to 0.5;
- the **random scaling size** σ_r , in a range from 0.0 to 0.1; and
- the **minibatch size** n_b , choosing from either 32, 64, 128, or 256;

The optimised hyperparameters found by SHERPA are then used to fully train each network for up to 30 epochs. Early stopping is used to stop training once the network specific performance metric has plateaued.

3.4.1 Training — cosmic rejection

The cosmic rejection network is trained on a sample of 3.15 million simulated events produced using the detector simulation and event generation methods outlined in section 2. Roughly $1/3^{rd}$ are ν_μ beam events, $1/3^{rd}$ ν_e beam events, and $1/3^{rd}$ cosmic muon events, the counts of which are shown in figure 4. All beam events (both ν_μ and ν_e) are generated using the expected unoscillated CHIPS-5 ν_μ energy spectrum to closely mimic the dominant ν_μ beam component and ν_e signal.




Figure 4. The number of training events per category for the cosmic rejection and beam classification network forms. In the cosmic rejection case, all beam event interaction types (shown in blue) are classed as beam events against the cosmic events (shown in red). In the beam classification case, only those events shown in blue are used.

The SHERPA optimised hyperparameters: $\eta_0 = 0.00005$, $c_d = 0.7$, $p_d = 0.1$, $\sigma_r = 0.02$, and $n_b = 128$ are used to train the network. The *cosmic score* accuracy is used for SHERPA optimisation and early stopping.

3.4.2 Training — beam classification

The beam classification network form is trained on a 1.67 million event subset of the training sample used for the cosmic rejection network form, excluding the cosmic muon events. The event counts are again shown in figure 4.

The SHERPA optimised hyperparameters: $\eta_0 = 0.0002$, $c_d = 0.5$, $p_d = 0.1$, $\sigma_r = 0.02$, and $n_b = 128$ are used to train the network. The *combined category* accuracy is used for SHERPA optimisation and early stopping.

3.4.3 Training — energy estimation

Accurate neutrino energy estimation is accomplished using multiple energy estimation networks trained on separate samples of ν_e and ν_μ events across all CC interaction types. This approach results in superior performance compared to a single energy estimation network or if just separate ν_e and ν_μ networks are trained. This is because a single set of network weights is unlikely to capture the specific topological features that contribute to the energy for all types of event.

Separate energy estimation networks are trained for each of CC-QEL (and CC-MEC), CC-Res, and CC-DIS for both ν_e and ν_μ events (6 in total) using 250000 simulated events each. Only events for which the primary charged lepton is fully contained within the detector volume are used for

training. Note that CC-QEL and CC-MEC energy estimation is combined into a single network as both have similar final state topologies (a single charged lepton).

The SHERPA optimised hyperparameters: $\eta_0 = 0.0002$, $c_d = 0.5$, $p_d = 0.1$, $\sigma_r = 0.0$, and $n_b = 128$ are used to train each network. The *neutrino energy* mean absolute error is used for SHERPA optimisation and early stopping.

4 Results

Although the time taken to train the CNNs can be extensive, once trained, the time required to calculate all network outputs (inference time) for a single event is on the order of 2 ms. When combined with event seeding and event map generation, the total time taken to fully reconstruct and classify a raw event is less than 0.1 seconds. When compared to the ~ 15 minutes required for each event using the standard reconstruction and classification methods, the difference is encouraging.

4.1 Evaluation sample

An independent sample of events is used to evaluate the combined performance of the trained CNNs. The evaluation sample consists of 400000 beam and 350000 cosmic muon simulated events. The beam events include the expected ν_μ , $\bar{\nu}_\mu$, ν_e and $\bar{\nu}_e$ components as well as events generated to mimic the ν_e signal component. During the evaluation, there is no discrimination made between neutrinos and anti-neutrinos.

Only the neutrino mode of the NuMI beam operation is considered here. CC neutrino events have higher hadronic energy fractions compared to antineutrino events. Therefore the efficiency for lepton identification could be lower on average for the neutrino events. Although this assumption was not tested here, Super-Kamiokande assumes similar efficiency for ν_μ and $\bar{\nu}_\mu$ events [51] but DUNE as an example of a CNN investigation, quotes higher $\bar{\nu}_\mu$ efficiency [52]. So the assumption that $\bar{\nu}_\mu$ will not be worse than ν_μ , is conservative.

All evaluation events are weighted to match the expected spectrum at the CHIPS-5 detector using the flux, cross-sections, and oscillation probabilities (derived from the NuFIT oscillation parameters [53], assuming the normal hierarchy, and including matter effects). Additional weighting also scales the sample to match data taking in the NuMI beam for a single year, corresponding to 6×10^{20} protons on target (POT). Cosmic muon events are weighted according to a 11.8 kHz expected CHIPS-5 rate over a full year. The final weighted spectrum of evaluation events is shown in figure 5.

4.2 Preselection cuts

A simple preselection is applied to all evaluation samples. Designed to reject cosmic and NC events while keeping the selection efficiency of CC beam events high, the preselection consists of four simple cuts, shown in figure 6. Firstly, the total number of collected photoelectrons (charge) across all PMTs in the event must be greater than 250. Secondly, the maximum Hough transform space height must be greater than 250 photoelectrons.¹ Thirdly, the seeding procedure $\cos(\theta)$ direction must be between ± 0.7 . Finally, the seeding procedure ϕ direction must be between ± 1.1 radians. The first two cuts reject low energy, NC events, while the last two reject cosmic events whose activity is not along the beam direction.

¹This is dependent on the particular 64 by 64 binning used for the Hough transform.




Figure 5. Weighted spectrum of events contained within the evaluation sample. The weighting is designed to mimic the expected event spectrum of the CHIPS-5 detector. Beam events are weighted by combining the expected unoscillated flux with cross-sections and standard oscillation probabilities, while cosmic events are weighted using the expected cosmic rate. Shown in blue, green, and olive are the surviving CC ν_μ , signal CC ν_e , and intrinsic beam CC ν_e spectra respectively, binned in terms of their neutrino energy. Shown in red is the NC event spectra, binned in terms of the energy of the hadronic component (excluding the outgoing neutrino energy) to represent more accurately the energy visible to the detector. Finally, shown in black is the cosmic muon event spectrum binned in terms of the muon energy.

4.3 Cosmic rejection and containment

The *cosmic score* output from the trained cosmic rejection network form shows an excellent separation between beam (output close to zero) and cosmic (output close to one) events, as can be seen in figure 7. Notably, 99.2% of beam events are associated with a score < 0.0001 . A *cosmic score* of below 0.0001 is therefore chosen to select beam-like events. Out of the total 350000 cosmic events in the evaluation sample, all are rejected by this cut after preselection.

For accurate neutrino energy estimation, the activity of an event must be fully contained within the volume of the detector. Within the CHIPS-5 detector, only 44% of the primary charged muons from CC ν_μ are fully contained within the detector volume. Therefore, the second output from the cosmic rejection network, *escapes score* is also used to select events. Although this output only considers the primary charged lepton, instead of all event particles, it still acts as a reasonable proxy for event containment. The distribution of *escapes score* output values for each event category is shown in figure 8.

An *escapes score* value below 0.33 is chosen to select events for which the primary charged lepton is deemed to be fully contained within the detector. This cut value is chosen to maximise the fraction of CC ν_μ events which are correctly classified as having their primary charged lepton contained or not, within the detector leading to $96.8 \pm 0.1\%$ of CC ν_μ events being classified correctly. As expected, the vast majority ($97.2 \pm 0.1\%$) of the short track CC ν_e and NC events are selected by this cut.

The total number of expected events per year that pass each successive cut (including preselection) for each event category is shown in table 1. Both CC ν_e categories are selected with an

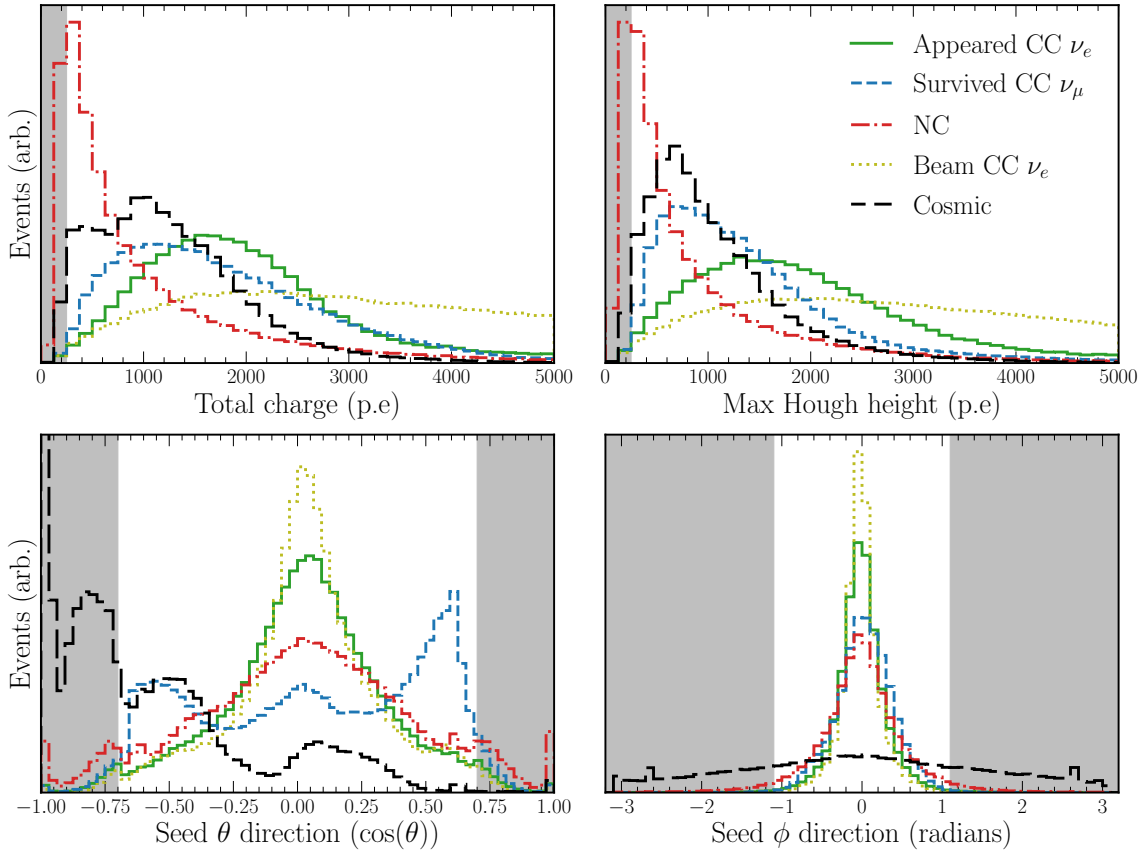


Figure 6. The four preselection cuts and their effect on the different event categories. The grey regions indicate rejected events.




Figure 7. Distribution of *cosmic score* output values from the trained cosmic rejection network for the different event categories. A score close to one signifies a cosmic-like event, while a score close to zero corresponds to a beam-like event. Only preselected events are shown to better highlight performance.




Figure 8. Distribution of *escapes score* output values from the trained cosmic rejection network for the different event categories. A score close to one signifies an uncontained event, while a score close to zero corresponds to a contained event. The containment cut value is shown at 0.33 with the arrow indicating the events that are selected.

Table 1. The total number of expected (weighted) events and the number that pass successive selection cuts for the different event categories. The preselection, *cosmic score* cut, and *escapes score* cut numbers are shown. The selection efficiency relative to the total number of events after all the cuts have been applied is also shown for each event category. For the cosmic events, an upper limit on the values is given.

Selection	App CC ν_e	CC ν_μ	Beam CC ν_e	NC	Cosmic
Total events	44.17 ± 0.02	2045.9 ± 33.3	35.06 ± 0.01	354.7 ± 5.8	2100000 ± 200000
+ Preselection	41.21 ± 0.02	1889.5 ± 30.9	33.52 ± 0.01	243.2 ± 4.0	430000 ± 40000
+ Cosmic cut	41.10 ± 0.02	1874.4 ± 30.7	33.35 ± 0.01	241.6 ± 4.0	$< 6 \pm 0.6$
+ Escapes cut	40.68 ± 0.02	795.7 ± 12.3	32.86 ± 0.01	233.0 ± 3.9	$< 6 \pm 0.6$
Cuts Eff	$92.1 \pm 0.1\%$	$38.9 \pm 0.1\%$	$93.7 \pm 0.1\%$	$65.7 \pm 0.3\%$	$< 2.9 \pm 0.3 \times 10^{-6}$

efficiency greater than 92%, relative to the total number of expected events, while CC ν_μ events have a $38.9 \pm 0.1\%$ selection efficiency, mainly due to the *escapes score* cut (as expected). NC events are found to be primarily rejected by the preselection, while cosmic events are rejected by both the preselection and *cosmic score* cuts.

An upper limit of $< 6 \pm 0.6$ cosmic muon events per year passing all the cuts is determined, showing that the cosmic muon rejection works well. Of all the true cosmic muon events with a *cosmic score* less than 0.9, none would be classified as signal CC ν_e events by the beam classification detailed in section 4.4. Therefore, the expected cosmic muon contamination of both beam selections (CC ν_e and CC ν_μ) is expected to be negligible relative to the selected number of signal events and is therefore ignored for the rest of this evaluation.

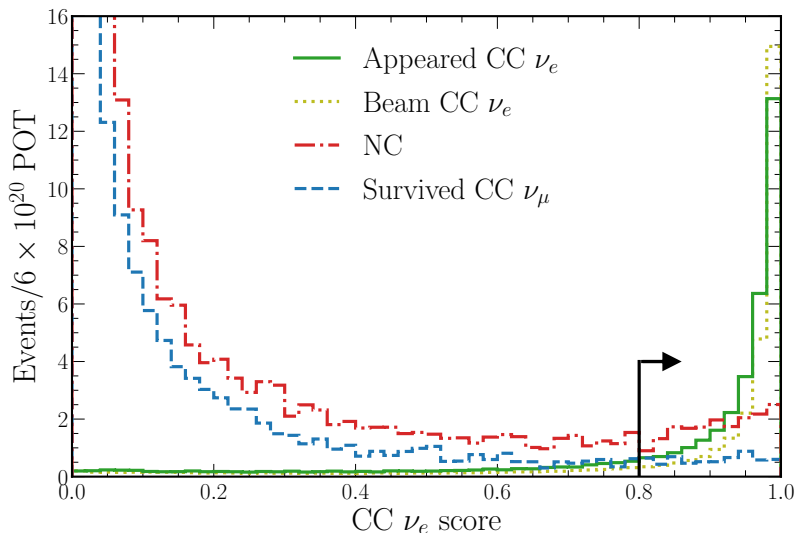


Figure 9. Distribution of *combined category* CC ν_e scores from the trained beam classification network for the different event categories. A score close to one signifies a CC ν_e like event. The FOM- ν_e optimised cut value is shown at 0.8 with the arrow indicating the events that are selected. The y-axis has been truncated so that the CC ν_μ and NC components are not fully visible to better show the distribution of signal CC ν_e events.

4.4 CC ν_e selection

The distribution of CC ν_e scores from the trained beam classification network for the different event categories are shown in figure 9. A strong separation between CC ν_e signal and both CC ν_μ and NC background events is achieved. As no attempt is made to separate the CC ν_e signal component from the intrinsic beam CC ν_e background, both are clustered with scores close to one as expected.

To assess the classification performance more rigorously, a selection score for each of the output categories found by maximising a figure-of-merit (FOM), is calculated. All events with a score above this optimised value are then deemed signal. To minimise the expected measurement statistical error, the value of efficiency \times purity (proportional to the square of $s/\sqrt{s+b}$) is optimised as the FOM [54].

The efficiency, purity, and their product (the FOM- ν_e) for CC ν_e events (both signal and beam background) as a function of selection score are shown in figure 10. The FOM- ν_e is optimised by selecting events with a CC ν_e score above 0.8, achieving a value of 0.519 ± 0.004 . Note that the FOM- ν_e is optimised considering both the signal and intrinsic beam CC ν_e components.

The total number of events is shown in table 2 with event category alongside the corresponding selection efficiencies and CC ν_e signal and combined CC ν_e purities. The purities are defined as the fraction of events within the selection that are true signal events. The final FOM- ν_e selected CC ν_e signal purity of $38.3 \pm 0.3\%$ may appear low, but this is mainly due to the indistinguishable intrinsic beam CC ν_e contamination. When both CC ν_e components are considered as signal, the selection purity is $70.9 \pm 0.6\%$.

The FOM- ν_e optimised CC ν_e selection efficiency, relative to the total number of events as a function of energy for the different event categories, is shown in figure 11 alongside the signal purity. From low neutrino energies, both CC ν_e category selection efficiencies rise to a plateau

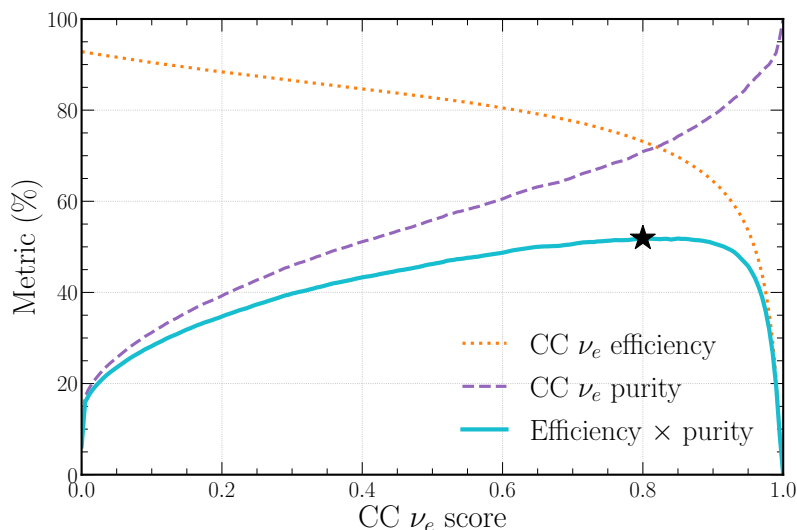


Figure 10. CC ν_e efficiency, purity, and efficiency \times purity (FOM- ν_e) curves for different values of CC ν_e score selection. The Maximum FOM- ν_e value of 0.519 ± 0.004 is indicated at a CC ν_e score of 0.8 by the star.

Table 2. Table showing CC ν_e selected event numbers and corresponding efficiencies for the various event categories as well as associated purities. Shown are the total event numbers, those after the preselection, *cosmic score* cut, and *escapes score* cut, in addition to the numbers after the FOM- ν_e optimised selection with the efficiencies relative to the total number of events shown for both. Both the signal CC ν_e purity and the joint signal and beam CC ν_e purity are shown for each selection.

Selection	CC ν_e sig	CC ν_μ bkg	CC ν_e bkg	NC bkg	Purity sig	Purity CC ν_e
Total events	44.17 ± 0.02	2045.9 ± 33.3	35.06 ± 0.01	354.7 ± 5.8	$1.78 \pm 0.02\%$	$3.19 \pm 0.03\%$
+ Cuts	40.68 ± 0.02	795.7 ± 12.3	32.86 ± 0.01	233.0 ± 3.9	$3.69 \pm 0.02\%$	$6.67 \pm 0.03\%$
+ FOM- ν_e	31.27 ± 0.02	6.0 ± 0.1	26.69 ± 0.01	17.8 ± 0.3	$38.3 \pm 0.3\%$	$70.9 \pm 0.6\%$
Cuts Eff	$92.1 \pm 0.1\%$	$38.9 \pm 0.1\%$	$93.7 \pm 0.1\%$	$65.7 \pm 0.3\%$	—	—
FOM- ν_e Eff	$70.8 \pm 0.2\%$	$0.29 \pm 0.02\%$	$76.1 \pm 0.1\%$	$5.0 \pm 0.2\%$	—	—

of approximately 80% beginning at 4 GeV. This is expected as low energy CC ν_e events have less well-defined electron Cherenkov rings, leading to their rejection. Problematically, this turn-on behaviour cuts into the true signal CC ν_e distribution, especially around the 1.5 GeV oscillation maximum. Future work should explore whether a CC ν_e selection cut that varies with energy can lead to a greater proportion of these low energy events being selected.

Due to the abundance of selected intrinsic beam CC ν_e events at higher energies, the signal CC ν_e purity is observed to peak at approximately 2.5 GeV (reasonably close to the oscillation maximum) before declining. Importantly, within the key signal region from 2 to 4 GeV, the signal CC ν_e purity is $> 55\%$, larger than the $38.3 \pm 0.3\%$ across the full evaluation sample. The NC efficiency is seen to slowly increase, approaching 15% for hadronic component energies above 5 GeV; this is due to misidentification of high energy pions or protons as electrons. However, within the key signal region, NC selection efficiency remains low.

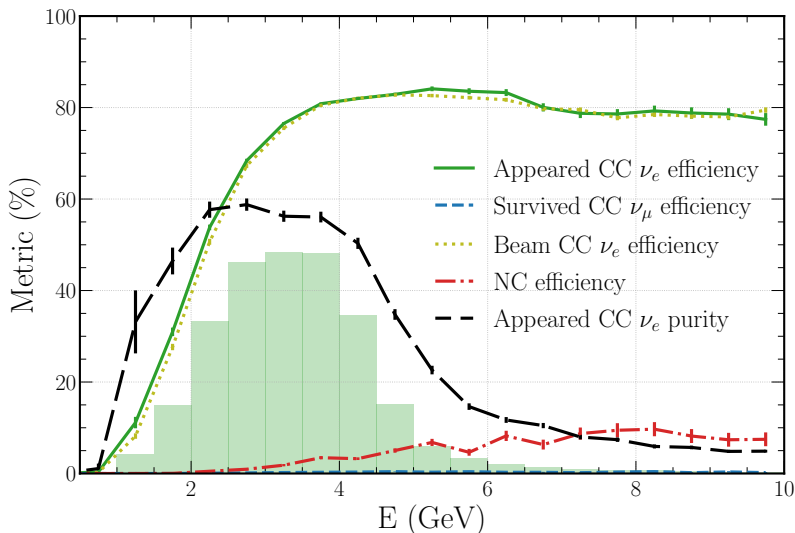


Figure 11. FOM- ν_e selection efficiencies relative to the total number of events for the different event categories as well as signal CC ν_e purity as a function of energy. All CC categories are shown in terms of the true neutrino energy, while NC events are shown in terms of the true hadronic component energy. The survived CC ν_μ efficiency is so low it is barely visible near zero. For reference, the true signal CC ν_e neutrino energy distribution is shown in green.

The best way to understand the relative performance of the CNN CC ν_e classification is by comparison with the standard event selection, outlined in detail in references [55] and [56]. A maximum efficiency \times purity of 0.132 ± 0.005 is achieved; only $\sim 25\%$ the value reached by the CNN approach. Both the combined signal and beam CC ν_e efficiency of $34.7 \pm 0.8\%$ compared to $73.4 \pm 0.2\%$ and purity of $39.3 \pm 1.2\%$ compared to $70.9 \pm 0.6\%$ are considerably lower than that provided by the new CNN classification.

Furthermore, the new signal CC ν_e signal efficiency of $70.8 \pm 0.2\%$ compares well to the 62% and 64% achieved by the NOvA and T2K CC ν_e selections, respectively. However, purity is significantly lower at $38.3 \pm 0.3\%$ compared to the 78% and 80% reached by NOvA and T2K [57, 58]. A large proportion of this difference can be explained by the lower neutrino energies and greater off-axis angles at which these experiments operate. Not only does this increase the proportion of easy to identify CC-QEL events, but it also reduces the indistinguishable beam CC ν_e contamination.

4.5 CC ν_μ selection

The distribution of CC ν_μ scores from the trained beam classification network form for the different event categories are shown in figure 12. Excellent separation between CC ν_μ signal and both CC ν_e components and NC background is achieved. For high CC ν_μ scores (close to one) the difference between signal and background rates is approximately three orders of magnitude.

The efficiency, purity, and their product (the FOM- ν_μ) for CC ν_μ events as a function of selecting events above a certain CC ν_μ score are shown in figure 13. FOM- ν_μ is optimised by selecting events with a CC ν_μ score above 0.315, achieving a value of 0.365 ± 0.002 . The total number of events, those selected by the previously mentioned cuts, and those additionally selected




Figure 12. Distribution of *combined category* CC ν_μ scores from the trained beam classification network for the different event categories. A score close to one signifies a CC ν_μ like event. The FOM- ν_μ optimised cut value is shown at 0.315 with the arrow indicating the events that are selected.

Table 3. Table showing CC ν_μ selected event numbers and corresponding efficiencies for the various event categories as well as the associated signal purity. Shown are the total event numbers, those after the preselection, *cosmic score* cut, and *escapes score* cut (Cuts), in addition to the numbers after the FOM- ν_μ optimised selection with the efficiencies relative to the total number of events shown for both. The CC ν_μ signal purity is also shown for each selection.

Selection	CC ν_μ sig	App CC ν_e bkg	Beam CC ν_e bkg	NC bkg	Purity sig
Total events	2045.9 \pm 33.3	44.17 \pm 0.02	35.06 \pm 0.01	354.7 \pm 5.8	82.5 \pm 0.2%
+ Cuts	795.7 \pm 12.3	40.68 \pm 0.02	32.86 \pm 0.01	233.0 \pm 3.9	72.2 \pm 0.2%
+ FOM- ν_μ	756.4 \pm 11.6	1.293 \pm 0.001	1.315 \pm 0.001	29.0 \pm 0.5	96.0 \pm 0.1%
Cuts Eff	38.9 \pm 0.1%	92.1 \pm 0.1%	93.7 \pm 0.1%	65.7 \pm 0.3%	—
FOM- ν_μ Eff	37.0 \pm 0.1%	2.9 \pm 0.1%	3.8 \pm 0.1%	8.2 \pm 0.2%	—

by the FOM- ν_μ optimised selection are shown in table 3 for each event category alongside the corresponding selection efficiencies and CC ν_μ signal purity.

The signal efficiency of 37.0 \pm 0.1% compares well to the 31% and 36% achieved by the NOvA and T2K CC ν_μ selections, respectively [57, 58]. This is also the case for the signal purity of 96.0 \pm 0.1% compared to the 98.6% and 94% purities of the NOvA and T2K selections. Although the final signal efficiency is low, this is desirable to ensure events are fully contained for energy estimation. When considering just those CC ν_μ events for which the primary charged muon is contained within the detector volume at the truth level, an 87.5 \pm 0.1% selection efficiency is achieved.

The FOM- ν_μ optimised CC ν_μ selection efficiency, relative to the total number of events as a function of energy for the different event categories, is shown in figure 14. The CC ν_μ selection efficiency peaks at just below 2 GeV before slowly declining as higher energy events are less likely




Figure 13. CC ν_μ efficiency, purity, and efficiency \times purity (FOM- ν_μ) curves for different values of CC ν_μ score selection. The maximum FOM- ν_μ value of 0.365 ± 0.002 is indicated at a CC ν_μ score of 0.315 by the star.




Figure 14. FOM- ν_μ selection efficiencies relative to the total number of events for the different event categories as well as the CC ν_μ purity as a function of energy. All CC categories are shown in terms of the true neutrino energy, while NC events are shown in terms of the true hadronic component energy. For reference, the true CC ν_μ neutrino energy distribution is shown in blue.

to have their primary charged muon fully contained within the detector. Of interest is the expected dip in the otherwise very high ($> 90\%$) CC ν_μ purity at approximately 1.5 GeV, corresponding to approximately the muon neutrino oscillation maximum where the signal to background ratio is low. As in the CC ν_e selection case, the NC efficiency is seen to rise with energy, again due to misidentification of energetic protons and pions.

Predicted category							
CC-Other	0.019	0.012	0.012	0.035	0.0037	0.16	0.016
CC-MEC	0	0	0	0	0	0	0
CC-Coh	3e-07	0.00032	0.00041	0.055	0	4e-06	0.00043
CC-DIS	0.018	0.16	0.8	0.29	0.00015	0.31	0.64
CC-Res	0.27	0.63	0.18	0.52	0.19	0.36	0.25
CC-QEL	0.69	0.19	0.014	0.11	0.81	0.17	0.097
	CC-QEL	CC-Res	CC-DIS	CC-Coh	CC-MEC	CC-Other	NC
	True category						

Figure 15. Classification matrix for the *CC category* output of the trained beam classification network. Shown are events that have either been selected by the CC ν_e or CC ν_μ selection. Events are simply classified using the categorical score for which they have the highest value. The numbers shown are the fraction of true category events classified into each of the six possible categories.

4.6 Interaction type classification

Using the *CC category* output of the trained beam classification network form, the CC interaction type for both CC ν_e and CC ν_μ selected events can be determined. As in the *combined category* output case, the highest-scoring neuron can be used for classification, resulting in the matrix shown in figure 15. Note that only events which are selected by either the CC ν_e or CC ν_μ selection are shown.

Reasonable classification accuracy greater than 60% is achieved across the three dominant interaction types, CC-QEL, CC-Res, and CC-DIS. The CC-Coh and CC-MEC types are found to be commonly misidentified as CC-Res and CC-QEL respectively, due to the imbalanced training dataset and their corresponding topological similarities. Background NC events which pass either CC ν_e or CC ν_μ selection (commonly high in energy) are found to be typically classified as CC-DIS; this is expected as they frequently contain multiple energetic particles in the final state.

4.7 Energy estimation

By using the CC classification described, the differences between CC interaction types can be exploited to improve neutrino energy, charged lepton energy, and interaction vertex position and time estimation. For events classified as either CC ν_e or CC ν_μ with an associated *CC category*

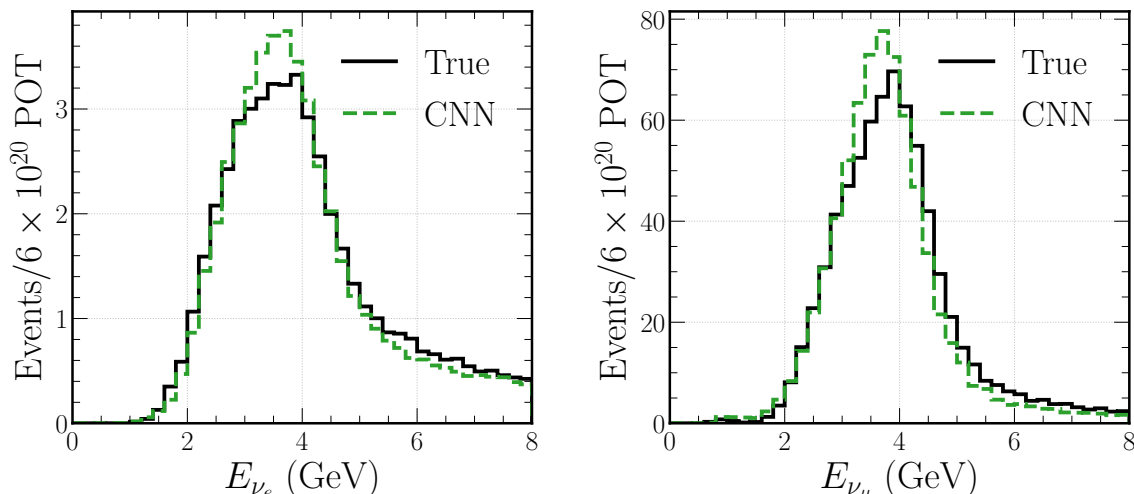


Figure 16. Distributions of true and CNN estimated neutrino energy for CC ν_e (left) and CC ν_μ (right) beam events. Only true CC ν_e and CC ν_μ events that are also selected by the CC ν_e and CC ν_μ selections respectively are shown.

Table 4. Summary of CC ν_e and CC ν_μ FWHM neutrino energy resolutions. Shown for each sample are the FWHM values for all selected signal events and the three dominant interaction type components, QEL, Res, and DIS. The FWHM values are calculated from the distributions shown in figure 17 and figure 18.

Event type	All	QEL component	Res component	DIS component
CC ν_e	$24.0 \pm 0.3\%$	$16.4 \pm 0.4\%$	$24.3 \pm 0.2\%$	$31.9 \pm 0.2\%$
CC ν_μ	$29.4 \pm 0.4\%$	$14.1 \pm 0.3\%$	$27.2 \pm 0.3\%$	$34.9 \pm 0.3\%$

interaction type, the corresponding bespoke energy estimation network outlined in section 3.4.3 is used for estimation.

Only three networks for each neutrino type are trained, one for each of the dominant interaction types CC-QEL (and CC-MEC), CC-Res, and CC-DIS. For events not classified by the *CC* category output as one of these categories, such as CC-Coh or CC-Other, the CC-Res network is used as it is the most topologically similar interaction type.

The distributions of CNN estimated *neutrino energy* output and true ν_e and ν_μ neutrino energies for true CC ν_e and CC ν_μ events respectively that are also selected by their corresponding CC selection are shown in figure 16. The CNN estimated distributions match the truth well across the full range of neutrino energies expected within CHIPS-5, except in the peak regions where the truth distribution shape is not fully captured.

To fully understand CNN energy estimation performance, histograms of ratios of fractional differences between CNN estimated (reco) and true neutrino energy for both true CC ν_e and CC ν_μ beam events that are also selected by their corresponding CC selection are shown in figure 17. Similar distributions splitting the signal components by interaction type are shown in figure 18. Furthermore, the Full Width Half Maximum (FWHM) neutrino energy resolutions derived from these plots for both CC ν_e and CC ν_μ events are shown in table 4.




Figure 17. Distributions of (reco-true)/true neutrino energies for both selected CC ν_e (left) and CC ν_μ (right) signal beam events.




Figure 18. Distributions of (reco-true)/true neutrino energies for both selected CC ν_e (left) and CC ν_μ (right) signal beam events by interaction type. The relative number of events between interaction types has been scaled for clearer comparison.

The interaction type FWHM values follow the expected pattern, with the single charged lepton QEL interactions achieving a smaller value than multi-particle DIS events. Furthermore, when the approximate resolution is derived from the FWHM values ($\text{FWHM}/(2\sqrt{2}\ln 2 \approx 2.355)$) the resolutions of $10.2 \pm 0.2\%$ and $12.5 \pm 0.2\%$ for CC ν_e and CC ν_μ respectively are comparable to the resolutions obtained by NOvA of 10.7% and 9.1% [57].

As in the CC ν_e selection case, the best way to understand the relative performance of the energy estimation is by comparison with the standard CHIPS reconstruction outlined in detail in references [55] and [56]. Although the standard reconstruction does not attempt to estimate the




Figure 19. Distributions of $(\text{reco-true})/\text{true}$ primary charged lepton energies for both CC ν_e (left) and CC ν_μ (right) beam events for both the new CNN approach and standard methods. Only QEL events are shown for clarity. The FWHM values of each distribution are indicated in the legends.

neutrino energy, the energy of the primary charged lepton in each CC event is predicted. The value can be compared to the *charged lepton energy* output of the energy estimation CNNs. Histograms of ratios of differences between CNN estimated (reco) and true charged lepton energy to true charged lepton energy for both CC ν_e and CC ν_μ beam QEL events are shown in figure 19.

A significant improvement is made using the new CNN approach. FWHM lepton energy improvements of 30% and 39% compared to the standard reconstruction values in QEL events for CC ν_e and CC ν_μ respectively are achieved, at $10.0 \pm 0.1\%$ and $9.0 \pm 0.1\%$, in their fractional percentage form. When the approximate resolution is found from the CC ν_e FWHM value ($4.2 \pm 0.1\%$), it compares well to the $\sim 2.5\%$ CC QEL charged lepton energy resolution reached by the Super-Kamiokande fitQun algorithm [59]. Given the significant differences in detector design, this is encouraging.

4.8 Comment

Recent CNN research has focused on the concern that they tend not to generalise well under distributional changes within the input data [60]. The CNN inputs can be characterised as being dependent on three sets of PMT information: the hit times, the hit charges, and the hit positions. The smearing of hit times, the smearing and shifting of hit charges, and the addition of random noise to both were used to test the robustness of the CNN to find that within reasonably expected ranges the performance change is negligible. During real world operation the simulation can also be tuned on real events to allow for anomalies between simulation and data to be accounted for and again the CNN made more robust.

Another common and justified concern with CNNs is their tendency to be used as a black box with little understanding of their inner working. For detailed physics analyses, this can have significant confidence implications for the final results. A technique to visualise trained CNNs is t-Distributed Stochastic Neighbour Embedding (t-SNE) [61]. The t-SNE procedure is an

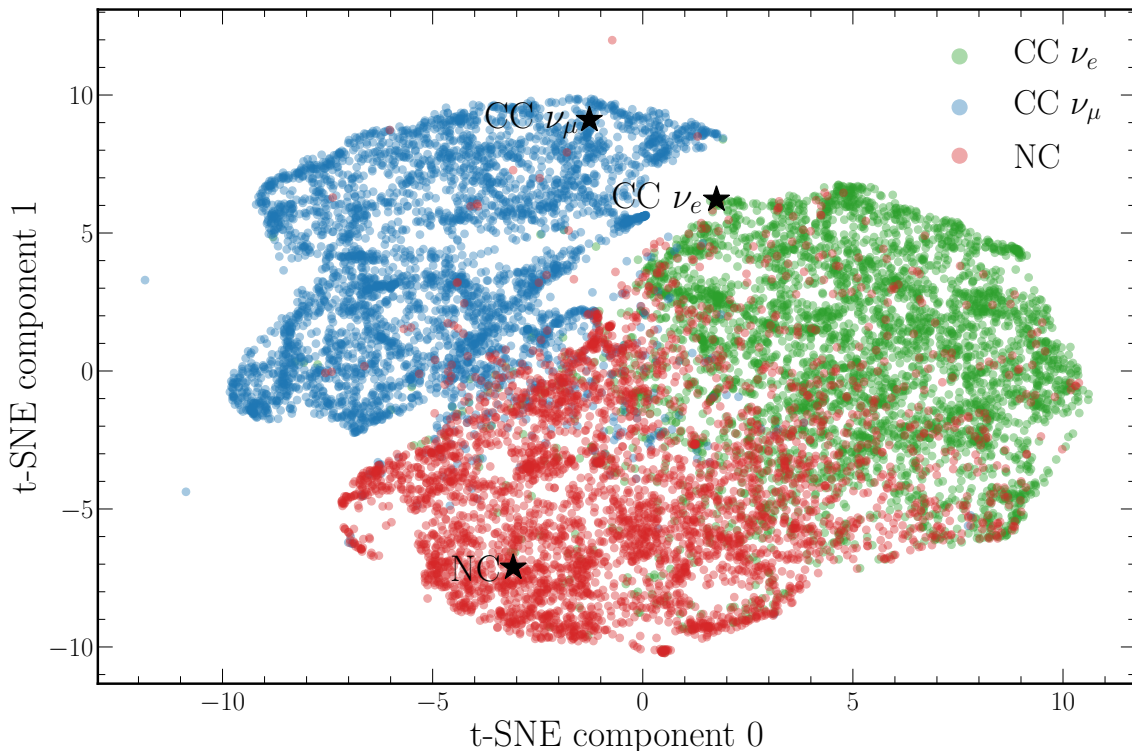


Figure 20. Two dimensional probability space of different beam events generated using the t-SNE procedure on the final fully-connected layer of the trained beam classification network. Three events, one highly CC ν_e like, one highly CC ν_μ like, and one highly NC like are highlighted and shown in figure 21.

unsupervised learning algorithm to visualise the learnt high-dimensional feature-space of a trained network in a lower number of dimensions. It accomplishes this by clustering events with similar features nearby in two-dimensional space and separating events with dissimilar features. Here, the outputs from the last fully connected layer before the output layer (with 512 dimensions) are used as input, as they provide the final representation of the learnt network features.

For the beam classification network form, three events, labelled in the t-SNE space of figure 20, are shown in figure 21. Each event is highly representative of its class, achieving a high respective *combined category* score. Both the CC ν_e and NC events are typical of that expected. However, the CC ν_μ event used here contains a muon that escapes the detector volume, this can be seen by the central peak in the hit-charge map. This topology suggests that strongly classified CC ν_μ events can be identified by this ‘escaping’ feature rather than the shape of the muon ring. Future work, therefore, should explore using only fully contained events during beam classification training.

5 Conclusions and future work

The CHIPS idea promotes a novel water Cherenkov based concept to counter the vast expense, increased complexity, and long construction time expected from future long-baseline neutrino oscillation experiments. It also makes use of modern machine learning techniques which have demonstrated a dramatic improvement in CHIPS-5 reconstruction and classification performance as presented here. Three forms of a baseline Convolutional Neural Network have been trained to reject

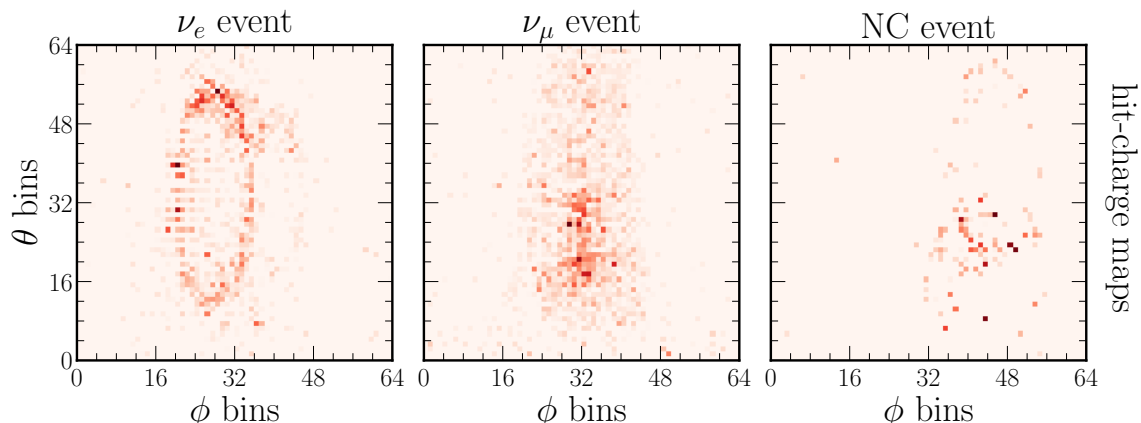


Figure 21. Hit-charge maps of the highly CC ν_e like, CC ν_μ like, and NC like events from figure 20.

Table 5. The key performance metrics for both the CC ν_e and CC ν_μ beam selections using the new CNN-based approach. The signal efficiency relative to the total number of expected events, the signal purity defined as the fraction of selected events which are signal, and the approximate signal neutrino energy resolution. The values considering both the appeared and beam CC ν_e components as signal are given for the CC ν_e selection.

Selection	Signal Efficiency	Signal Purity	$\sim \nu$ Energy Resolution
CC ν_e	$73.4 \pm 0.2\%$	$70.9 \pm 0.6\%$	$10.2 \pm 0.2\%$
CC ν_μ	$37.0 \pm 0.1\%$	$96.0 \pm 0.1\%$	$12.5 \pm 0.2\%$

cosmic muon events, classify beam events, and estimate neutrino energies, all using only the raw detector event as input. This new approach replaces the standard likelihood-based reconstruction and simple neural network classification, whilst greatly increasing generalisability and processing speed.

The new CNN-based approach is found to provide excellent performance in selecting an efficient and pure CC ν_e signal sample for which the neutrino energy can be accurately determined. In some cases, the performance is comparable with similar experiments, which is encouraging given the significant differences in detector design. Without the help of a veto the rejection of the 11.8 kHz cosmic muon background is still found to be strong. Less than $2.9 \pm 0.3 \times 10^{-6}$ of all event is accepted, equivalent to $< 6 \pm 0.6$ cosmic muon events contaminating the beam sample per year, of which none are expected to be classified as CC ν_e events.

The key performance metrics for both the CC ν_e and CC ν_μ beam selections are summarised in table 5. The learnt representation of the inputs found by the CNNs is seen to have strong discriminating power between categories when visualised using the t-SNE technique.

Although the results presented in this work are compelling, there are still clear avenues for exploration and improvement principally related to the critical performance drivers outlined.

Firstly, generating the input event maps to focus on the underlying Cherenkov profiles is important; therefore, methodology to further remove distortions, or more accurately determine the interaction vertex position will be beneficial. Secondly, the distribution of events (in energy or type) used within the training sample heavily impacts performance; thus, a comprehensive study of this behaviour could optimise the sample used. Finally, multi-task learning clearly shows promise, with further trial-and-error or a more generalised approach likely to uncover additional valuable tasks.

Acknowledgments

This work has been supported by European Research Council grant n° 694920 — CHROMIUM, The Department of Physics and Astronomy at University College London and the Department of Physics at the University of Wisconsin Madison.

References

- [1] P. Adamson et al., *CHerenkov detectors In mine PitS (CHIPS) letter of intent to FNAL*, in the proceedings of the *Snowmass 2013: Snowmass on the Mississippi*, (2013) [[arXiv:1307.5918](#)] [[INSPIRE](#)].
- [2] P. Adamson et al., *The NuMI neutrino beam*, *Nucl. Instrum. Meth. A* **806** (2016) 279 [[arXiv:1507.06690](#)] [[INSPIRE](#)].
- [3] F. Psihas, M. Groh, C. Tunnell and K. Warburton, *A review on machine learning for neutrino experiments*, *Int. J. Mod. Phys. A* **35** (2020) 2043005 [[arXiv:2008.01242](#)] [[INSPIRE](#)].
- [4] K. Fukushima, *Neocognitron: a self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position*, *Biol. Cyber.* **36** (1980) 193.
- [5] I. Goodfellow, Y. Bengio and A. Courville, *Deep learning*, MIT Press, U.S.A. (2016).
- [6] A. Aurisano et al., *A convolutional neural network neutrino event classifier*, *2016 JINST* **11** P09001 [[arXiv:1604.01444](#)] [[INSPIRE](#)].
- [7] C. Szegedy et al., *Going deeper with convolutions*, in the proceedings of the 2015 *IEEE conference on computer vision and pattern recognition (CVPR)*, (2015) [[DOI:10.1109/cvpr.2015.7298594](#)].
- [8] F. Psihas et al., *Context-enriched identification of particles with a convolutional network for neutrino events*, *Phys. Rev. D* **100** (2019) 073005 [[arXiv:1906.00713](#)] [[INSPIRE](#)].
- [9] P. Baldi, J. Bian, L. Hertel and L. Li, *Improved energy reconstruction in NOvA with regression convolutional neural networks*, *Phys. Rev. D* **99** (2019) 012011 [[arXiv:1811.04557](#)] [[INSPIRE](#)].
- [10] ICECUBE collaboration, *Application of deep neural networks to event type classification in IceCube*, *PoS ICRC2019* (2020) 937 [[arXiv:1908.08763](#)] [[INSPIRE](#)].
- [11] E. Racah et al., *Revealing fundamental physics from the Daya Bay neutrino experiment using deep neural networks*, in the proceedings of the 2016 15th *IEEE international conference on machine learning and applications (ICMLA)*, (2016) [[DOI:10.1109/icmla.2016.0160](#)].
- [12] KM3NET collaboration, *Event reconstruction for KM3NeT/ORCA using convolutional neural networks*, *2020 JINST* **15** P10005 [[arXiv:2004.08254](#)] [[INSPIRE](#)].
- [13] MICROBooNE collaboration, *Design and construction of the MicroBooNE detector*, *2017 JINST* **12** P02017 [[arXiv:1612.05824](#)] [[INSPIRE](#)].
- [14] MICROBooNE collaboration, *Convolutional neural network for multiple particle identification in the MicroBooNE liquid argon time projection chamber*, *Phys. Rev. D* **103** (2021) 092003 [[arXiv:2010.08653](#)] [[INSPIRE](#)].
- [15] MICROBooNE collaboration, *Convolutional neural networks applied to neutrino events in a liquid argon time projection chamber*, *2017 JINST* **12** P03011 [[arXiv:1611.05531](#)] [[INSPIRE](#)].
- [16] DUNE collaboration, *Neutrino interaction classification with a convolutional neural network in the DUNE far detector*, *Phys. Rev. D* **102** (2020) 092003 [[arXiv:2006.15052](#)] [[INSPIRE](#)].

- [17] DUNE collaboration, *Deep Underground Neutrino Experiment (DUNE), far detector technical design report, volume II: DUNE physics*, [arXiv:2002.03005](#) [INSPIRE].
- [18] DUNE collaboration, *Deep-learning-based kinematic reconstruction for DUNE*, [arXiv:2012.06181](#) [INSPIRE].
- [19] M. Rossi and S. Vallecorsa, *Deep learning strategies for ProtoDUNE raw data denoising*, *Comput. Softw. Big Sci.* **6** (2022) 2 [[arXiv:2103.01596](#)] [INSPIRE].
- [20] B. Graham, M. Engelcke and L. van der Maaten, *3D semantic segmentation with submanifold sparse convolutional networks*, [arXiv:1711.10275](#) [INSPIRE].
- [21] DEEPLearnPhysics collaboration, *Scalable deep convolutional neural networks for sparse, locally dense liquid argon time projection chamber data*, *Phys. Rev. D* **102** (2020) 012005 [[arXiv:1903.05663](#)] [INSPIRE].
- [22] K. Carloni, N.W. Kamp, A. Schneider and J.M. Conrad, *Convolutional neural networks for shower energy prediction in liquid argon time projection chambers*, *2022 JINST* **17** P02022 [[arXiv:2110.10766](#)] [INSPIRE].
- [23] J. Tingey et al., *Chips WCSIM git repository*, <https://gitlab.com/chipsneutrino/chips-sim> (2020).
- [24] *WCSim git repository*, <https://github.com/WCSim/WCSim> (2021).
- [25] GEANT4 collaboration, *GEANT4 — a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250 [INSPIRE].
- [26] J. Allison et al., *GEANT4 developments and applications*, *IEEE Trans. Nucl. Sci.* **53** (2006) 270 [INSPIRE].
- [27] J. Allison et al., *Recent developments in GEANT4*, *Nucl. Instrum. Meth. A* **835** (2016) 186 [INSPIRE].
- [28] DUNE collaboration, *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): conceptual design report, volume 1: the LBNF and DUNE projects*, [arXiv:1601.05471](#) [INSPIRE].
- [29] M. Campbell, *Measuring neutrino oscillations in the NOvA and CHIPS detectors*, Ph.D. thesis, University College, London, U.K. (2020).
- [30] *Photomultiplier tubes — basics and applications*, https://www.hamamatsu.com/resources/pdf/etd/PMT_handbook_v3aE.pdf (2020).
- [31] MINERvA collaboration, *Neutrino flux predictions for the NuMI beam*, *Phys. Rev. D* **94** (2016) 092005 [*Addendum ibid.* **95** (2017) 039903] [[arXiv:1607.00704](#)] [INSPIRE].
- [32] C. Andreopoulos et al., *The GENIE neutrino Monte Carlo generator*, *Nucl. Instrum. Meth. A* **614** (2010) 87 [[arXiv:0905.2517](#)] [INSPIRE].
- [33] C. Andreopoulos et al., *The GENIE neutrino Monte Carlo generator: physics and user manual*, [arXiv:1510.05494](#) [INSPIRE].
- [34] C. Haggmann et al., *Cosmic-ray shower library (CRY)*, https://nuclear.llnl.gov/simulation/doc_cry_v1.7/cry.pdf (2012).
- [35] C.A. Haggmann, D.J. Lange and D.M. Wright, *Monte Carlo simulation of proton-induced cosmic ray cascades in the atmosphere*, https://nuclear.llnl.gov/simulation/doc_cry_v1.7/cry_physics.pdf, Lawrence Livermore National Lab (LLNL), Livermore, CA, U.S.A. (2007).

- [36] S.I. Klimushin, E.V. Bugaev and I.A. Sokalski, *Precise parametrizations of muon energy losses in water*, in the proceedings of the 27th international cosmic ray conference, (2001) [[hep-ph/0106010](#)] [[INSPIRE](#)].
- [37] *WCSim: how to add your own photodetector tutorial*, <https://github.com/WCSim/WCSim/wiki/Tutorial:-How-to-add-your-own-photodetector> (2020).
- [38] J. Tingey et al., *Chips reconstruction git repository*, <https://gitlab.com/chipsneutrino/chips-reco> (2020).
- [39] L. Alzubaidi et al., *Review of deep learning: concepts, CNN architectures, challenges, applications, future directions*, *J. Big Data* **8** (2021) 1.
- [40] J. Tingey et al., *chipsnet git repository*, <https://gitlab.com/chipsneutrino/chips-net> (2020).
- [41] M. Abadi et al., *TensorFlow: large-scale machine learning on heterogeneous distributed systems*, [arXiv:1603.04467](#) [[INSPIRE](#)].
- [42] T. Theodore, *Particle identification in Cherenkov detectors using convolutional neural networks*, M.Sc. thesis, University of Toronto, Toronto, ON, Canada (2016).
- [43] P. Mukhopadhyay and B.B. Chaudhuri, *A survey of Hough transform*, *Pattern Recognit.* **48** (2015) 993.
- [44] K. Simonyan and A. Zisserman, *Very deep convolutional networks for large-scale image recognition*, [arXiv:1409.1556](#).
- [45] S. Ioffe and C. Szegedy, *Batch normalization: accelerating deep network training by reducing internal covariate shift*, in *Proceedings of the 32nd international conference on international conference on machine learning — volume 37, ICML'15*, (2015), p. 448 [[arXiv:1502.03167](#)] [[INSPIRE](#)].
- [46] J. Hu et al., *Squeeze-and-excitation networks*, *IEEE Trans. Pattern Anal. Machine Intell.* **42** (2020) 2011.
- [47] F. Chollet et al., *Keras*, <https://keras.io>.
- [48] R. Cipolla, Y. Gal and A. Kendall, *Multi-task learning using uncertainty to weigh losses for scene geometry and semantics*, in the proceedings of the 2018 *IEEE/CVF conference on computer vision and pattern recognition*, (2018) [[DOI:10.1109/cvpr.2018.00781](#)].
- [49] D.P. Kingma and J. Ba, *Adam: a method for stochastic optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [50] L. Hertel et al., *Sherpa: robust hyperparameter optimization for machine learning*, *SoftwareX* **12** (2020) 100591.
- [51] SUPER-KAMIOKANDE collaboration, *Sensitivity of Super-Kamiokande with Gadolinium to low energy anti-neutrinos from pre-supernova emission*, *Astrophys. J.* **885** (2019) 133 [[arXiv:1908.07551](#)] [[INSPIRE](#)].
- [52] DUNE collaboration, *Neutrino interaction classification with a convolutional neural network in the DUNE far detector*, *Phys. Rev. D* **102** (2020) 092003 [[arXiv:2006.15052](#)] [[INSPIRE](#)].
- [53] I. Esteban et al., *The fate of hints: updated global analysis of three-flavor neutrino oscillations*, *JHEP* **09** (2020) 178 [[arXiv:2007.14792](#)] [[INSPIRE](#)].
- [54] B. List, *Why and when to optimize efficiency times purity*, ETH Zurich internal note, ETH, Zurich, Switzerland (2002).
- [55] J. Tingey, *Convolutional neural networks for the CHIPS neutrino detector R&D project*, Ph.D. thesis, University College, London, U.K. (2021).

- [56] A. Blake et al., *CHIPS event reconstruction and design optimisation*, [arXiv:1612.04604](#) [INSPIRE].
- [57] NOvA collaboration, *First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA*, *Phys. Rev. Lett.* **123** (2019) 151803 [[arXiv:1906.04907](#)] [INSPIRE].
- [58] T2K collaboration, *Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6×10^{20} protons on target*, *Phys. Rev. D* **91** (2015) 072010 [[arXiv:1502.01550](#)] [INSPIRE].
- [59] SUPER-KAMIOKANDE collaboration, *Atmospheric neutrino oscillation analysis with improved event reconstruction in Super-Kamiokande IV*, *PTEP* **2019** (2019) 053F01 [[arXiv:1901.03230](#)] [INSPIRE].
- [60] J. Djolonga et al., *On robustness and transferability of convolutional neural networks*, [arXiv:2007.08558](#).
- [61] L. Maaten and G. Hinton, *Visualizing data using t-SNE*, *J. Mach. Learn. Res.* **9** (2008) 2579.