



# Maximum Likelihood Reconstruction of Water Cherenkov Events With Deep Generative Neural Networks

Mo Jia<sup>1</sup>, Karan Kumar<sup>1</sup>, Liam S. Mackey<sup>2</sup>, Alexander Putra<sup>3</sup>, Cristovao Vilela<sup>4\*</sup>, Michael J. Wilking<sup>1</sup>, Junjie Xia<sup>5</sup>, Chiaki Yanagisawa<sup>1,6\*</sup> and Karan Yang<sup>7</sup>

<sup>1</sup> Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY, United States, <sup>2</sup> Department of Physics, Applied Physics and Astronomy, Rensselaer Polytechnic Institute, Troy, NY, United States, <sup>3</sup> Department of Mathematics, BMCC/City University of New York, New York, NY, United States, <sup>4</sup> CERN European Organization for Nuclear Research, Geneva, Switzerland, <sup>5</sup> Institute for Cosmic Ray Research, Research Center for Cosmic Neutrinos, University of Tokyo, Chiba, Japan, <sup>6</sup> Department of Science, BMCC/City University of New York, New York, NY, United States, <sup>7</sup> Information Science, Cornell Tech, New York, NY, United States

## OPEN ACCESS

### Edited by:

Bo Jayatilaka,  
Fermi National Accelerator Laboratory  
(DOE), United States

### Reviewed by:

Gabriel Nathan Perdue,  
Fermilab Accelerator Complex, Fermi  
National Accelerator Laboratory  
(DOE), United States  
Anushree Ghosh,  
University of Padua, Italy

### \*Correspondence:

Cristovao Vilela  
c.vilela@cern.ch  
Chiaki Yanagisawa  
chiaki.yanagisawa@stonybrook.edu

### Specialty section:

This article was submitted to  
Big Data and AI in High Energy  
Physics,  
a section of the journal  
Frontiers in Big Data

Received: 02 February 2022

Accepted: 04 May 2022

Published: 17 June 2022

### Citation:

Jia M, Kumar K, Mackey LS, Putra A,  
Vilela C, Wilking MJ, Xia J,  
Yanagisawa C and Yang K (2022)  
Maximum Likelihood Reconstruction  
of Water Cherenkov Events With Deep  
Generative Neural Networks.  
Front. Big Data 5:868333.  
doi: 10.3389/fdata.2022.868333

Large water Cherenkov detectors have shaped our current knowledge of neutrino physics and nucleon decay, and will continue to do so in the foreseeable future. These highly capable detectors allow for directional and topological, as well as calorimetric information to be extracted from signals on their photosensors. The current state-of-the-art approach to water Cherenkov reconstruction relies on maximum-likelihood estimation, with several simplifying assumptions employed to make the problem tractable. In this paper, we describe neural networks that produce probability density functions for the signals at each photosensor, given a set of inputs that characterizes a particle in the detector. The neural networks we propose allow for likelihood-based approaches to event reconstruction with significantly fewer assumptions compared to traditional methods, and are thus expected to improve on the current performance of water Cherenkov detectors.

**Keywords:** experimental particle physics, event reconstruction, water Cherenkov detectors, generative models, convolutional neural network

## 1. INTRODUCTION

In high energy physics, several large water Cherenkov detectors have been used since 1980s such as IMB (Irvine-Michigan-Brookhaven) (Becker-Szendy et al., 1993), Kamiokande (Kamioka Nucleon Decay) (Oyama et al., 1989), Super-Kamiokande (SK) (Fukuda et al., 2003), and for the near future Hyper-Kamiokande (HK) (Abe et al., 2018) and proposed THEIA (Askin et al., 2020) and ESSnUSB (Alekou et al., 2021) experiments. These are a type of detector that uses Cherenkov radiation produced by charged particles traveling faster than the speed of light in water. Photons in radiation traverse on a conical surface with its axis in the direction of parent charged particle. Photomultiplier tubes (PMTs) mounted on the walls of the detector detect these photons. PMTs produce electric charges proportional to the number of photons detected. The amount of charges and the arrival times are digitally recorded by electronic circuits. Depending on the direction of the parent charged particle with respect to the surface instrumented with PMTs, a detected Cherenkov ring leaves a pattern in the shape of a circle, an ellipse, or a parabola. For ultrarelativistic particles in water, Cherenkov photons are emitted at an angle of approximately  $42^\circ$  with respect to the particle direction. This angle decreases as particles slow down due to energy losses in the water.

A good illustration of an atmospheric neutrino detected by the SK detector is available on the SK official site<sup>1</sup>. Water Cherenkov detectors contributed to the first detection of neutrinos from a supernova (Hirata et al., 1987), the discovery of neutrino oscillation in atmospheric neutrinos (Fukuda et al., 1998a), the confirmation of solar neutrino oscillation (Fukuda et al., 1998b) and search for proton decays (Takenaka et al., 2020).

Neutrino and nucleon decay experiments pose a particular set of event reconstruction challenges. In order to overcome the smallness of neutrino cross sections and long nucleon lifetimes, these detectors are designed to have as much active mass as possible and, unlike typical collider or fixed-target experiments, the location of the events within the detector is not known, even approximately, *a priori*. This challenge is compounded by a rich neutrino-nucleus interaction phenomenology (Alvarez-Ruso et al., 2018) at the energies of interest for many such experiments, in the order of GeV, which makes the detailed reconstruction of the events' final-state topology crucial both for precision measurements and potential discoveries.

Machine learning (ML) techniques have been increasingly adopted (Radovic et al., 2018) to tackle these challenges across several detector technologies, from segmented scintillator detectors (Aurisano et al., 2016; Perdue et al., 2018; Alonso-Monsalve et al., 2021) to liquid argon time-projection chambers (Abi et al., 2020; Abratenko et al., 2021). Most of these efforts have focused on the *discriminative* aspect of ML, with algorithms designed to classify events, for example into signal and background categories, or to estimate a variable of interest, such as the energy of an interacting neutrino.

In this work, we investigate a complementary approach by developing a ML-based *generative* model that encodes the likelihood for the measurements at each PMT as a function of variables describing the event. This likelihood function can then be used to reconstruct events, for example using gradient-descent methods to find the event hypothesis that maximizes the likelihood, or by sampling the likelihood with Markov Chain Monte Carlo methods. This approach emphasizes interpretability by allowing for the detailed examination of the likelihood surface for each event, while guaranteeing powerful discrimination between competing event hypotheses by virtue of the Neyman-Pearson lemma (Neyman and Pearson, 1933).

While we have no knowledge of other uses of generative models for event reconstruction tasks, similar models have been developed as a computationally efficient substitute for event simulation in collider experiments (Paganini et al., 2018; Vallecorsa et al., 2019; Butter and Plehn, 2020; Alanazi et al., 2021).

## 2. WATER CHERENKOV EVENT RECONSTRUCTION

An event in a water Cherenkov detector consists of a set of charges and times recorded by each PMT. PMTs for which the amount of charge collected exceeds a given threshold will

produce a hit, with a respective charge and time. An event is often broken into different hit clusters in time (subevents) and most read-out electronics systems effectively limit each PMT to measure a single charge (integrated over a period of time) and time (usually the time at which the electronic signal crosses the hit threshold) per subevent. While these detectors can in principle have arbitrary shapes, most currently running and proposed experiments are cylindrical, with an instrumented barrel, top and bottom end-caps. Our work so far focuses on such cylindrical geometries, though it can in principle be adapted to other detector shapes.

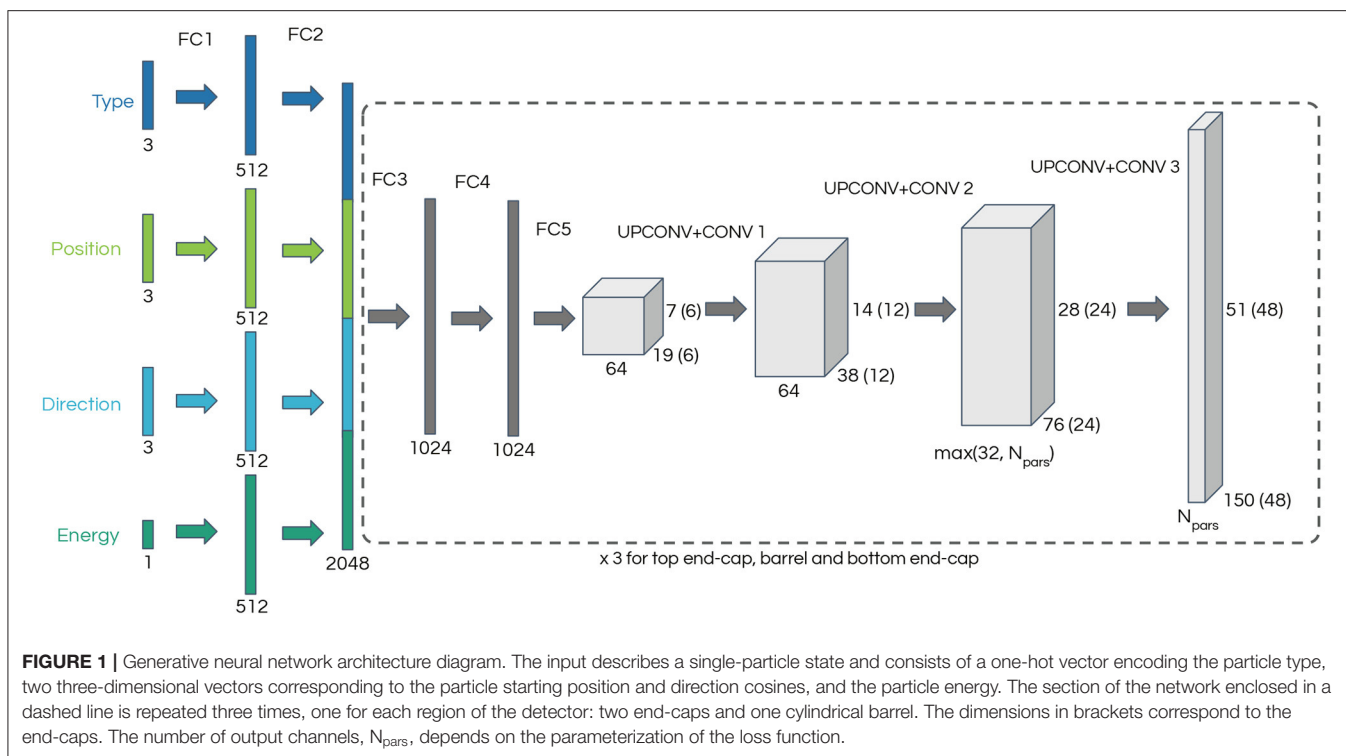
The state-of-the-art in water Cherenkov event reconstruction is the FiTQun maximum likelihood estimation algorithm, whose adoption has led to improved neutrino oscillation measurements by long-running experiments (Jiang et al., 2019; Abe et al., 2020). At the core of FiTQun is a likelihood function which is evaluated over every PMT in the detector, regardless of whether or not it registers a hit in the event:

$$L(x) = \prod_j^{unhit} P_j(unhit|x) \prod_i^{hit} \{1 - P_i(unhit|x)\} f_q(q_i|x) f_t(t_i|x), \quad (1)$$

where  $x$  denotes a set of parameters describing one or more particles in the detector, namely their type, starting positions, directions and momenta. The index  $j$  runs over all the PMTs that did not register a hit and the index  $i$  runs over all the hit PMTs, with  $P_j(unhit|x)$  being the probability of a certain PMT not registering a hit under the  $x$  hypothesis. The probability density function (hereafter PDF) for the observed charge  $q_i$  in the  $i^{\text{th}}$  hit PMT is  $f_q(q_i|x)$ , while the PDF for the observed time  $t_i$  for the same PMT is  $f_t(t_i|x)$ .

Event reconstruction proceeds by minimizing the negative log-likelihood  $-\ln L(x)$  by using the MINUIT gradient-descent package (James, 1994), with the hypothesis  $x$  which minimizes  $-\ln L(x)$  taken as the best-fit hypothesis for the event. This process is repeated for elements of  $x$  which are categorical in nature, such as the particle type, or the number of particles (and their types) in the event. Likelihood ratio tests are used to discriminate between competing categorical hypotheses. In order to make this likelihood function tractable, it is factorized into several low-dimension components. In particular, the PDFs associated to Cherenkov photons that produce a hit without having scattered in the water or reflected in the detector surfaces are factorized from the PDFs that describe so-called indirect photons that scatter or reflect before producing a PMT hit. The level of detail of the indirect photon PDF is limited by the high number of dimensions required for it to be fully specified. In particular, this limitation makes it difficult to reconstruct heavier, typically slower, particles such as protons, since it is challenging to accommodate the effect of the decreasing Cherenkov photon emission angle. Finally, each component of the factorized likelihood needs to be tuned separately to the detector geometry of interest, requiring a large amount of bespoke simulated data with different components of the simulation disabled in turn.

<sup>1</sup><http://www-sk.icrr.u-tokyo.ac.jp/sk/detector/cherenkov-e.html>



### 3. GENERATIVE NEURAL NETWORKS FOR MAXIMUM LIKELIHOOD RECONSTRUCTION

In order to overcome the challenges in water Cherenkov event reconstruction associated to the curse of dimensionality, we have designed a convolutional generative neural network to replace the factorized likelihood function in FiTQun. The outputs of this network are PDFs for the hit charges and times at each PMT in the detector. Given the success of ML methods in processing high dimensional data, we do not factorize the likelihood function in our approach, nor do we require bespoke sets of simulated data for training. Rather, the neural network can be trained using a regular, fully detailed, simulated data set.

Like in the existing FiTQun algorithm, the likelihood function for multiple particles in an event can in principle be combined to form complex event hypotheses. While this capability is a future goal of our project, in the present work we have focused on demonstrating the method using single-particle events consisting of either a showering electron, or a track-like muon.

#### 3.1. Network Architectures

We have implemented our model using the *PyTorch* (Paszke et al., 2019) framework, and we used the network architecture in Dosovitskiy et al. (2014) as a starting point. As depicted in **Figure 1**, our network is composed of two parts: five fully connected (FC) layers are followed by three pairs of transposed convolution (UPCONV) and convolution (CONV) layers. We have found that pairing the two types of convolutional layers as done in Dosovitskiy et al. (2014) results in smoother outputs.

However, as discussed below this approach may also limit the network's ability to reproduce sharp features in our data. Rectified linear unit (*ReLU*) activation, defined as  $f(x) = 0$  for  $x < 0$  and  $f(x) = x$  elsewhere, is used after each layer.

The neural network input is a vector describing a single-particle state, consisting of a three-dimensional one-hot encoding of the particle type (electron, muon, or gamma<sup>2</sup>), a three-dimensional vector with the particle starting position, a three-dimensional vector of the particle direction cosines and finally a scalar corresponding to the particle energy. Following the architecture in Dosovitskiy et al. (2014), each input type is processed through two fully connected layers of 512 nodes each on its own (FC1 and FC2). The output is concatenated into a feature vector of length 2,048 which passes through two fully connected layers of 1024 nodes (FC3 and FC4) followed by a final fully connected layer (FC5) which results in the starting point for the convolutional part of the network. The three convolutional layers result in images where each pixel represents a PMT in the detector and the pixel values encode parameters which are used to build the likelihood function.

The portion of the network starting with FC3 is repeated three times, one for each section of the cylindrical detector: the top and bottom end-caps, and the cylindrical barrel. Each of these detector regions is represented as two-dimensional images with  $48 \times 48$  pixels for the end-caps and  $150 \times 51$  pixels for the barrel. Only the FC1 and FC2 layers are common to all neural network segments, with the weights of FC3, FC4,

<sup>2</sup>We do not train with gammas in this iteration, but instead include a place holder for it.

FC5 and the convolutional layers being independent for each section of the detector. The depth of the output depends on the parameterization of the likelihood function and ranges from four in the simplest case to 61 in the most complex case.

A kernel size of  $4 \times 4$  and a stride of  $2 \times 2$  is used for all UPCONV layers and the CONV layers use a  $3 \times 3$  kernel and  $1 \times 1$  stride. The first UPCONV + CONV pair layer takes a tensor of depth 64 and produces an output of depth of 64, and the second layer reduces the depth to either 32 or the number of output channels, whichever is largest. The final layer produces the desired number of output channels, which depends on the choice of loss function. Padding is used on the output to match the odd dimension of the detector barrel.

To match the square images produced by the neural network to the circular end-caps of the detector, the pixels close to the corners, which do not correspond to physical PMTs, are masked when evaluating the loss function.

### 3.2. Loss Functions

We have designed the loss function in a similar fashion to Equation (1), with two components: one describing the probability of the PMT being hit and the other describing the probability density function for a hit charge and time:

$$\text{Loss} = -\log L = \sum_i -\log P_{\text{unhit}}(y_i) + \sum_{i_{\text{hit}}} -\log p_{qt}(q_{i_{\text{hit}}}, t_{i_{\text{hit}}}), \tag{2}$$

where the index  $i$  runs over all the PMTs in the detector,  $y_i$  is a label set to 1 if the PMT is not hit or 0 if the PMT is hit,  $i_{\text{hit}}$  runs only over the PMTs which are hit in the event, and  $p_{qt}(q_{i_{\text{hit}}}, t_{i_{\text{hit}}})$  is the PDF for observing charge  $q_{i_{\text{hit}}}$  and time  $t_{i_{\text{hit}}}$ . The loss function is the sum of the negative log-likelihood over all PMTs in the three regions of the detector.

For the PMT hit probability, we use the *PyTorch* built-in function `BCEWithLogitsLoss` which implements the binary cross-entropy loss (Cover and Thomas, 1991) (equivalent to the negative log-likelihood) using a *Sigmoid* function to regularize predictions of hit probability. A single channel of the neural network output represents the logit of the hit probability.

It was observed in Xia et al. (2019) and this work that both the charge and timing of a PMT’s responses can be very non-Gaussian, despite the stochastic process of photoelectron multiplication, due to a list of reasons ranging from the kinematics of event to the reflection of light from the ambient environment. In order to accommodate the *a priori* unknown functional form of the hit charge and time PDF, we approximate this function with a weighted mixture of Gaussian PDFs in one dimension (taking into account only the hit charge) or two dimensions (hit charge and time). **Figure 2** shows an example of the charge PDF predicted by networks with different numbers of Gaussian components for an arbitrarily chosen PMT.

We have explored several combinations of the number of components in the Gaussian mixture in one and two dimensions, and different parameterizations for the two-dimensional Gaussian function, either keeping the charge and time uncorrelated, or including a correlation factor. Each component in the Gaussian mixture is weighted by a coefficient

which corresponds to one of the neural network output channels. To preserve the normalization of the PDFs, a *Softmax* function is applied to the coefficients.

In the case where the hit charge and time are treated as uncorrelated in each of the Gaussian components, the mixture of Gaussians PDF is given by:

$$-\log p_{qt}(q_{i_{\text{hit}}}, t_{i_{\text{hit}}}) = -\sum_{i_{\text{hit}}} \left[ \sum_j \left( \log(n_j) - \log(\sqrt{2\pi}\sigma_{q_j}) - \frac{(q_{i_{\text{hit}}} - \mu_{q_j})^2}{2\sigma_{q_j}^2} - \log(\sqrt{2\pi}\sigma_{t_j}) - \frac{(t_{i_{\text{hit}}} - \mu_{t_j})^2}{2\sigma_{t_j}^2} \right) \right]_{i_{\text{hit}}}, \tag{3}$$

where  $i_{\text{hit}}$  runs over the hit PMTs, as in Equation (2),  $j$  runs over the  $N$  Gaussian components,  $n_j$  is the normalization factor for the  $j$ -th component,  $\mu_{q_j}$  and  $\mu_{t_j}$  are the charge and time means for the  $j$ -th component, respectively, and  $\sigma_{q_j}$  and  $\sigma_{t_j}$  are the corresponding standard deviations. For the one-dimensional case, where only the hit charges are considered but not the times, the third line of the equation is omitted. In order to improve numerical stability, *PyTorch*’s implementation of the log-sum-exp function is used to evaluate the loss function.

For each component in the Gaussian mixture, a set of network output channels correspond to  $\log \mu_{q_j}$ ,  $\log \sigma_{q_j}$ ,  $\mu_{t_j}$ , and  $\log \sigma_{t_j}$ , with the logarithms being used to ensure non-negative values of the hit charge and standard deviations. Together with the hit probability and the Gaussian component coefficients, the total number of channels is  $1 + N \times 3$  for the one-dimensional case and  $1 + N \times 5$  for the two-dimensional case.

In the bivariate case, correlations between the charge and time dimensions of the loss function above can arise from the independent sets of  $\mu_{q_j}$  and  $\mu_{t_j}$ . However, the most general form of the two-dimensional mixture of Gaussians includes a covariance term describing the correlation between charge and time within each of the mixture’s components. To realize this goal we formulate two-dimensional PDF of correlated charge and time responses in the way introduced by Williams (1996):

$$f(\boldsymbol{\eta} | \boldsymbol{\theta}) = \sum_j \frac{n_j}{(2\pi)|\Sigma_j|^{1/2}} \exp\left(-\frac{1}{2}(\boldsymbol{\eta} - \boldsymbol{\theta}_j)^T \Sigma_j^{-1}(\boldsymbol{\eta} - \boldsymbol{\theta}_j)\right), \tag{4}$$

where  $\boldsymbol{\eta}$  and  $\boldsymbol{\theta}$  include both the charge and time as a two-dimensional vector:

$$\boldsymbol{\eta} = (t, q) \tag{5}$$

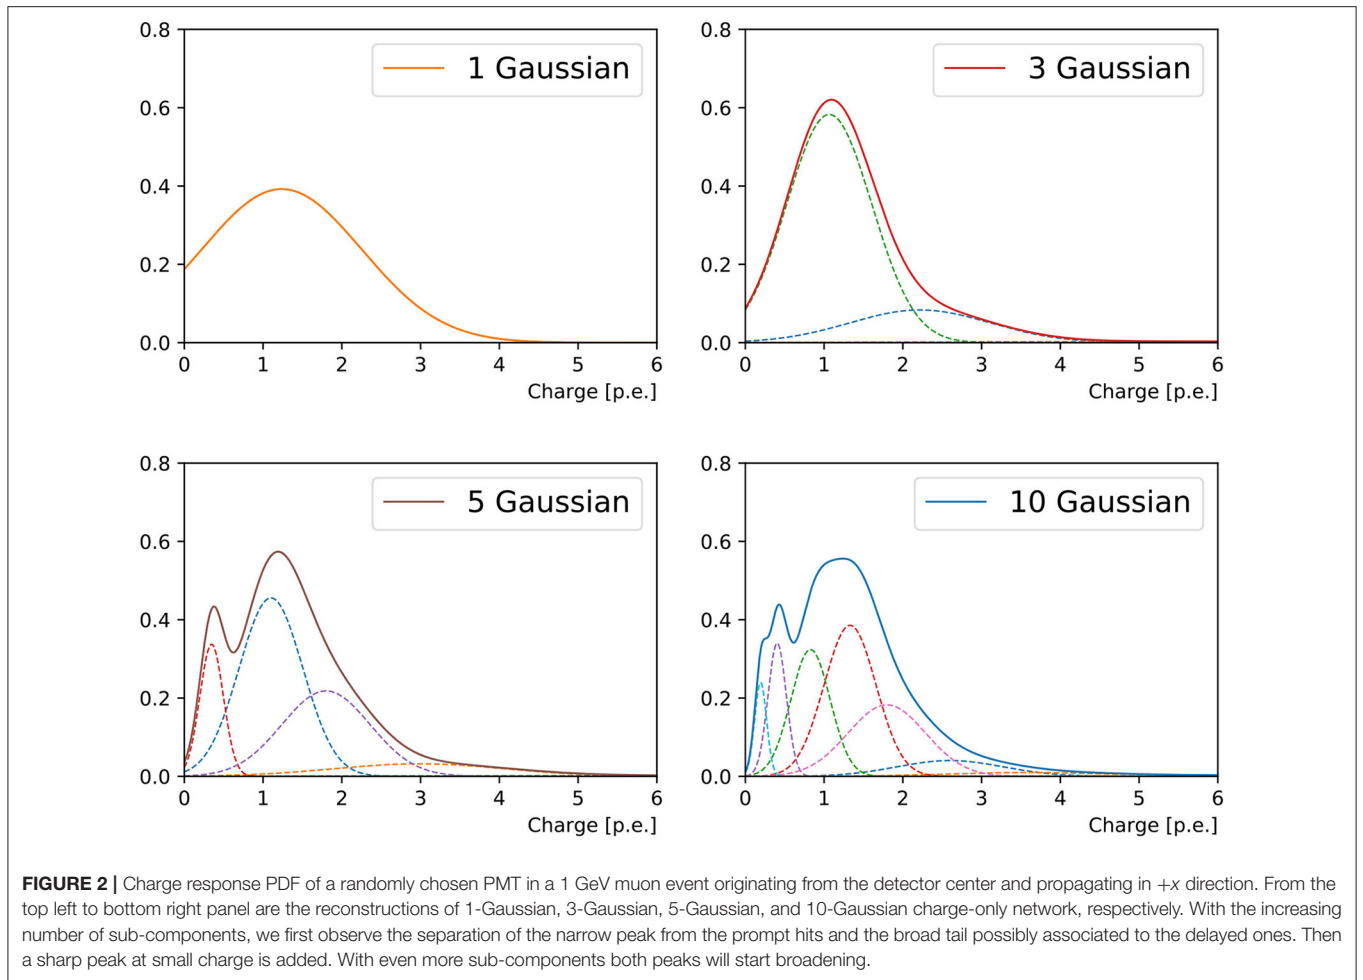
$$\boldsymbol{\theta}_j = (\mu_t, \mu_q)_j \tag{6}$$

where  $\Sigma_j$  is a two-dimensional covariance matrix.

To improve numerical stability we use triangular matrix instead of the full covariance matrix  $\Sigma$  by Cholesky decomposition:

$$\Sigma^{-1} = \begin{pmatrix} \alpha_{11} & 0 \\ \alpha_{12} & \alpha_{22} \end{pmatrix} \begin{pmatrix} \alpha_{11} & \alpha_{12} \\ 0 & \alpha_{22} \end{pmatrix} \tag{7}$$





with positive-only diagonal terms and

$$|\Sigma|^{-1/2} = \alpha_{11}\alpha_{22} \tag{8}$$

The  $\alpha$ 's are also predicted by the neural network, requiring a total of  $1 + 6 \times N$  output channels. To ensure positive definite matrix  $\Sigma$ , we keep the absolute value of  $\alpha_{12}$  while forcing its sign so that  $(\eta - \theta_i)^T \Sigma_i^{-1} (\eta - \theta_i) > 0$  for  $\forall \eta, \theta$ .

$$\alpha_{12} \begin{cases} < 0, & \text{if } (t - \mu_t)(q - \mu_q) < 0 \\ > 0, & \text{if } (t - \mu_t)(q - \mu_q) > 0 \end{cases}$$

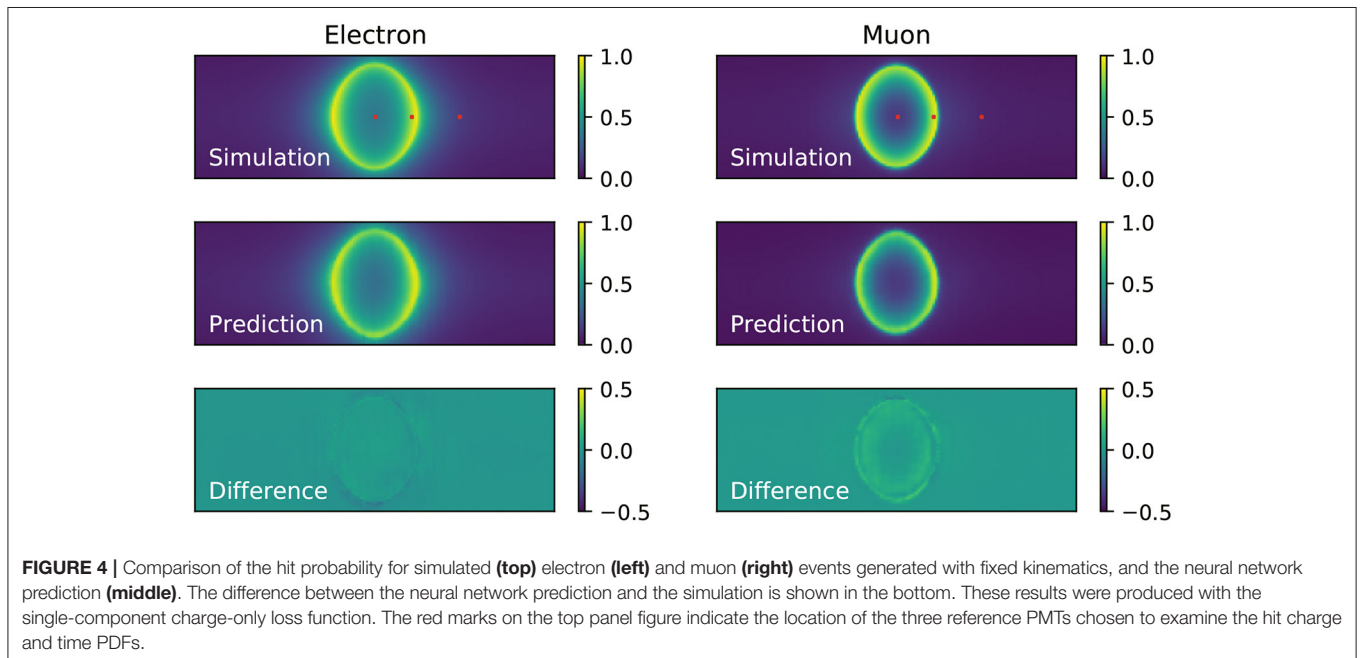
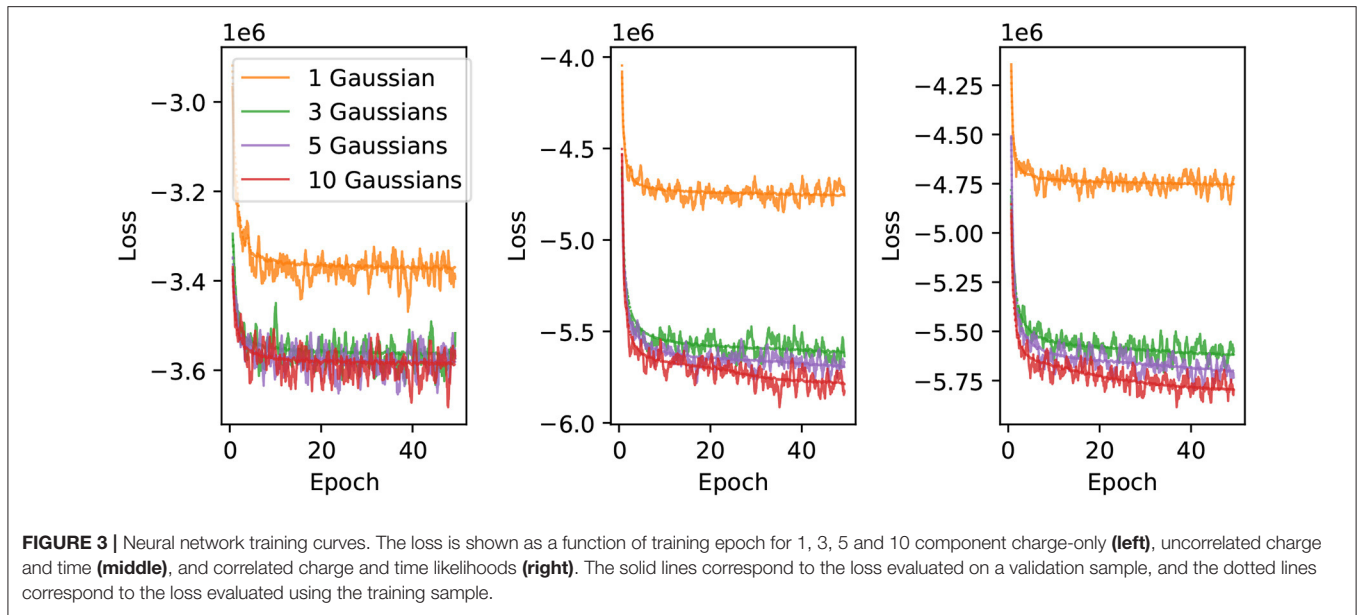
#### 4. DATA SETS

The data set used in this work consists of single-electron and single-muon events generated with random kinematics and positions in a model of the SK detector—a cylindrical volume with 36 meters in height and 34 meters in diameter. The coordinate system used to describe the detector has  $z$  pointing along the cylinder axis. The data are simulated with the water

Cherenkov detector simulation package WCSim<sup>3</sup>, an open-source program based on Geant4 (Agostinelli et al., 2003) and ROOT (Brun and Rademakers, 1997), which models the particle propagation in the detector and the electronic response of the PMTs. For each simulated event, its Monte Carlo truth, including particle type, position, direction cosines and energy, is saved in order to be used during neural network training. The hit charges and times at every PMT in each event are stored in three two-dimensional arrays representing the unrolled cylindrical barrel of the detector ( $151 \times 50$ ) and the two circular end-caps ( $48 \times 48$ ). The data set consists of 1,003,200 electron events and 1,049,415 muon events. The kinetic energy of the electrons is uniformly distributed between 1 and 6,500 MeV and muons between 150 and 6,500 MeV. Spontaneous discharges in the PMTs, a source of uncorrelated noise commonly referred to as the dark rate, are not simulated, for simplicity. Taking into account the PMTs' dark rate is a future goal of this project. In order to produce pure single-particle data sets, delayed electrons originating from muon decays are not simulated.

In addition to the training data set, we have also produced another set of muon and electron events with fixed position,

<sup>3</sup><https://github.com/WCSim/WCSim>

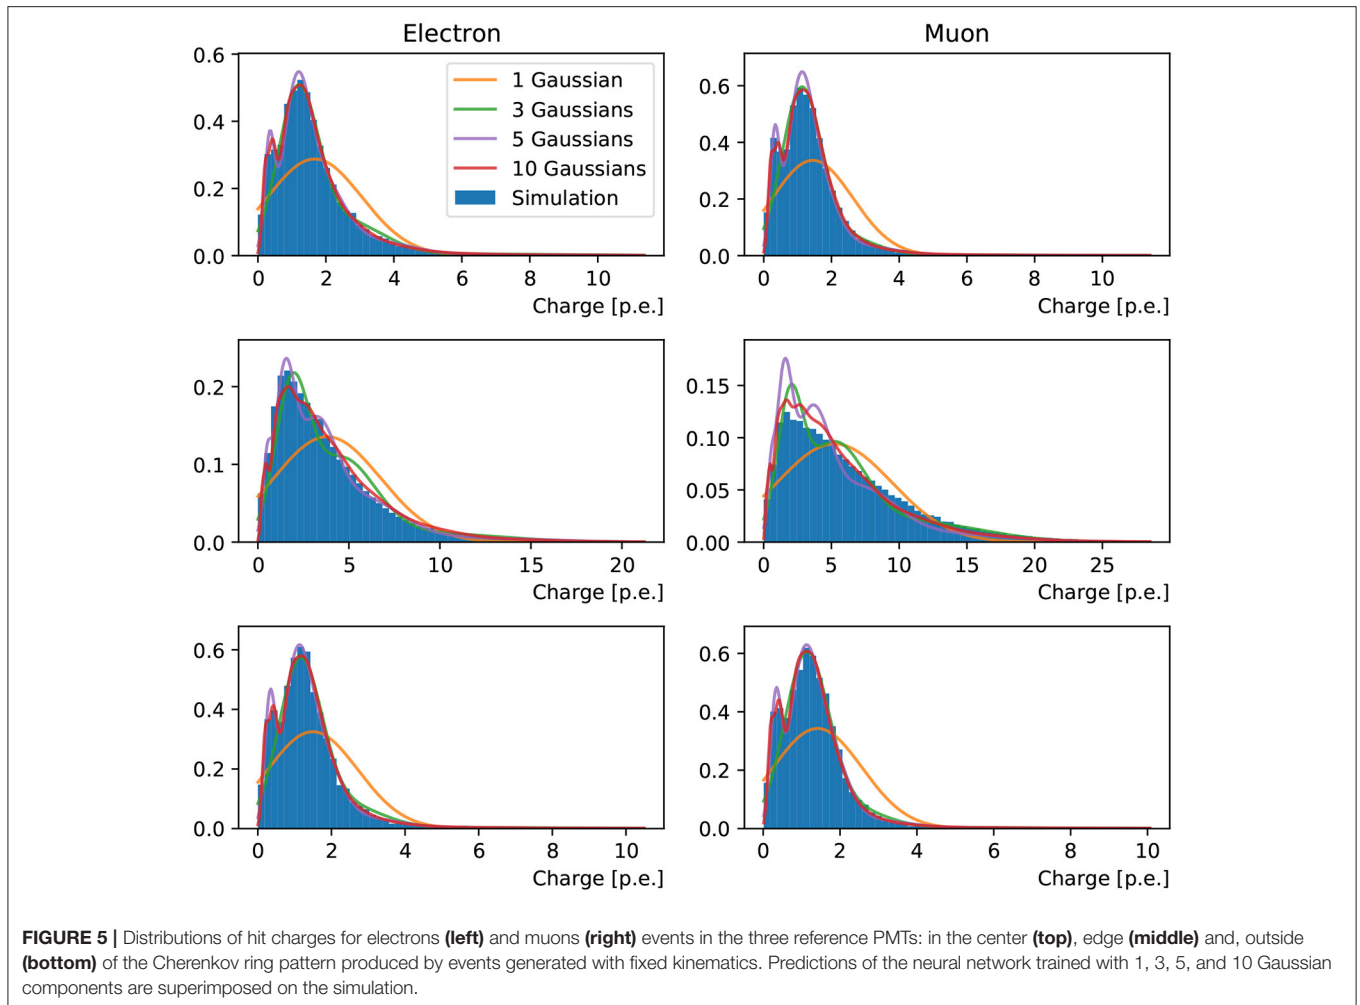


direction, and energy. These events originate in the center of the detector and the particles propagate along the  $x$  direction (onto the cylinder barrel) with a kinetic energy of 500 MeV. These events are used to evaluate the quality of the likelihood functions generated by the neural networks by comparing them to the hit probability extracted directly from the simulation as well as the distributions of simulated hit charges and times.

## 5. TRAINING

We use 75% of the randomized electron and muon events for training, with the remaining 25% used for *in-situ* validation

every 100 iterations. The network is trained using minibatches of 200 shuffled events. The results presented in this work were obtained by training the neural networks for 50 epochs, using the Adam (Kingma and Ba, 2017) optimizer with the initial learning rate set to 0.0002, and all other parameters left at the *PyTorch* default values. To improve numerical stability we normalize the particle position by the detector dimensions, the energy by 5,000 MeV, PMT hit charges by 2,500 p.e., and convert PMT hit times to  $\mu\text{s}$  with an offset of  $-1 \mu\text{s}$  so that the values are of  $\mathcal{O}(10^{-3} \sim 1)$ . We found that this pre-normalization gave the training more stability compared to using BatchNorm layers on the inputs, and using both BatchNorm layers and input pre-normalization resulted in slower training. **Figure 3**



shows examples of training curves for different configurations of the neural network. The configurations with a single Gaussian component or one-dimensional loss function generally converge within 50 epochs, while networks with multiple components and two-dimensional loss functions still show downward trends at the end of 50 epochs of training. Therefore, the results presented in this publication can likely be improved with extended training of the networks. We have saved both the trained network weights and optimizer states at 50 epochs, which can be used for further training beyond this point.

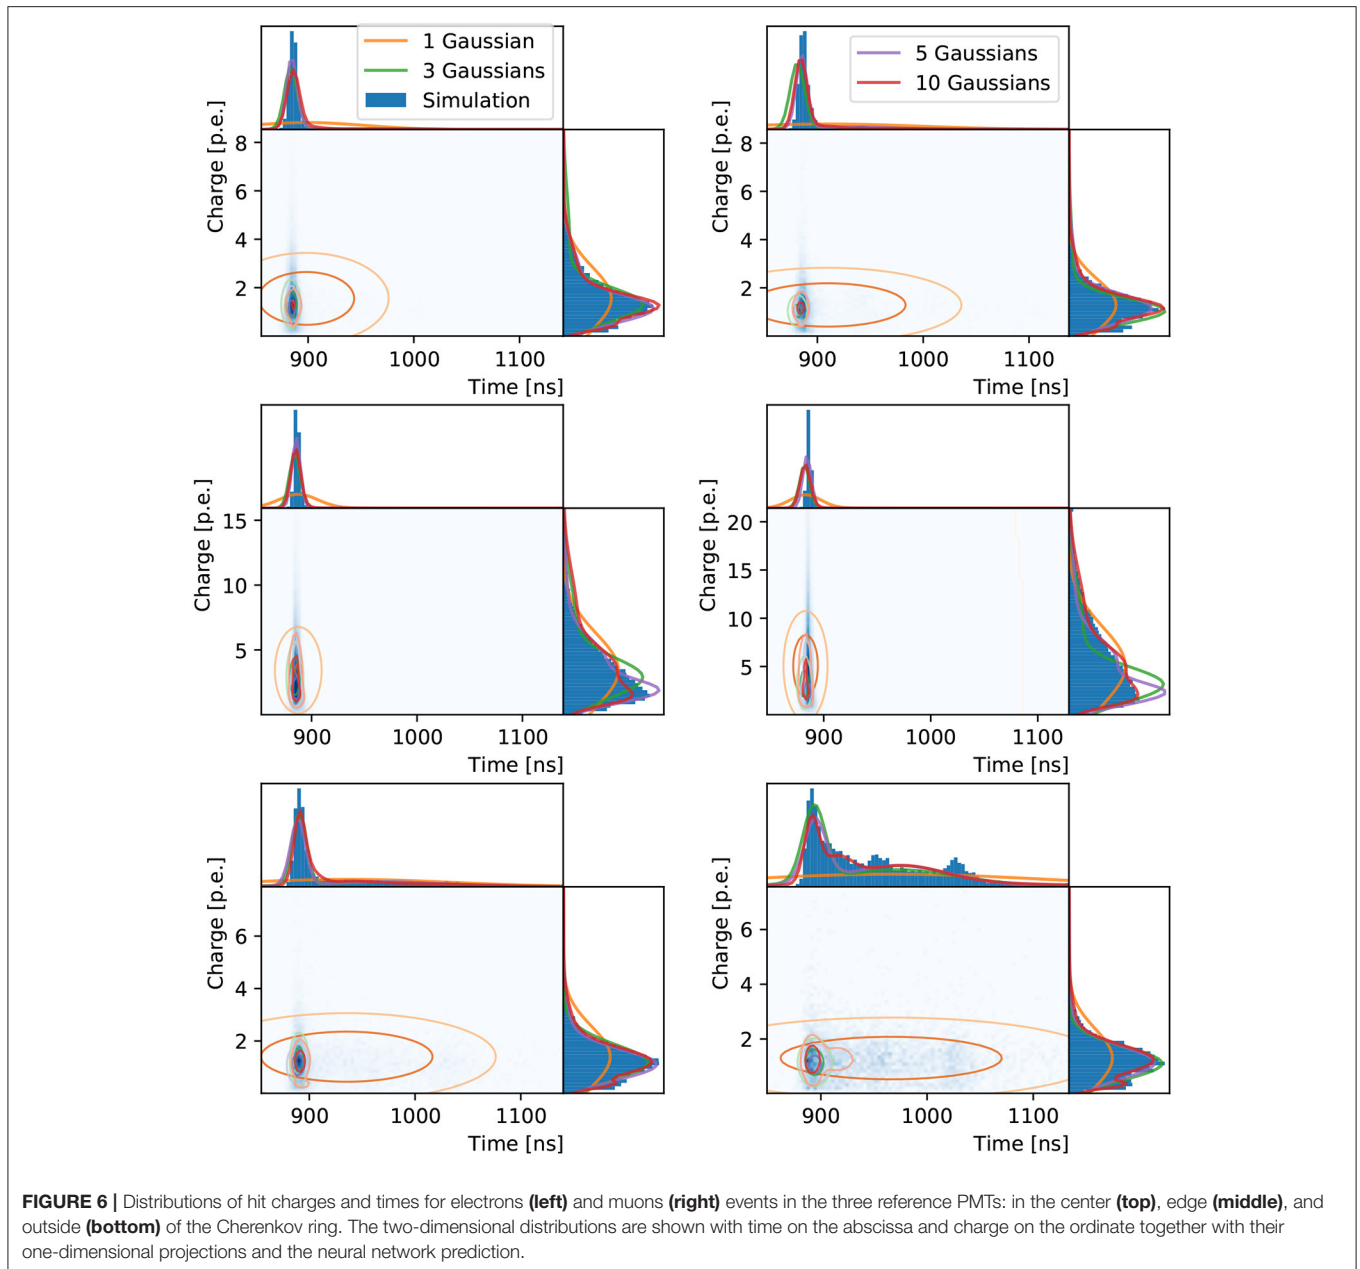
## 6. RESULTS

For the neural networks described in this work to be effective in maximum likelihood reconstruction, they need to describe the data as closely as possible, to avoid bias, and they should be smooth to avoid local minima in the likelihood surface which would make event reconstruction challenging. In this section we examine the output of the neural network in two ways: first we compare the neural network prediction to the distributions obtained directly from the simulated data set produced with fixed

particle kinematics; we then use the randomized validation data set to scan the likelihood function as a function of the particle energy to evaluate the smoothness, bias, and the accuracy of the neural network in identifying the true particle type in the events.

### 6.1. Comparison of Neural Network Output to Events Simulated With Fixed Kinematics

To evaluate the accuracy of the neural network at predicting the probability of PMTs in the detector being hit, we measure the hit probability in the simulation by counting the number of times each PMT is hit in the set of events with fixed kinematics and dividing each PMT's hit count by the total number of events in the set. The resulting hit probabilities are shown in **Figure 4**, where the Cherenkov ring pattern is clearly seen and the difference between the fuzzy rings produced by showering electrons and sharp rings produced by track-like muons is evident. For the electron events, the agreement between the neural network prediction (middle panel of the same figure) and the simulation is excellent. On the other hand, for muon events there is some level of residual difference (bottom panel) that might indicate the neural network might have limited ability

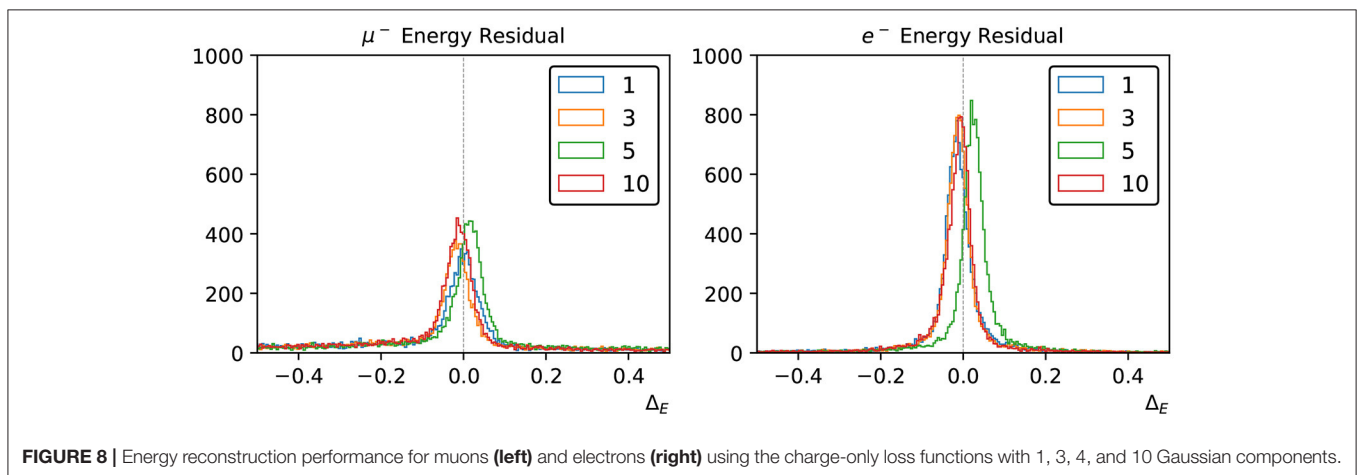
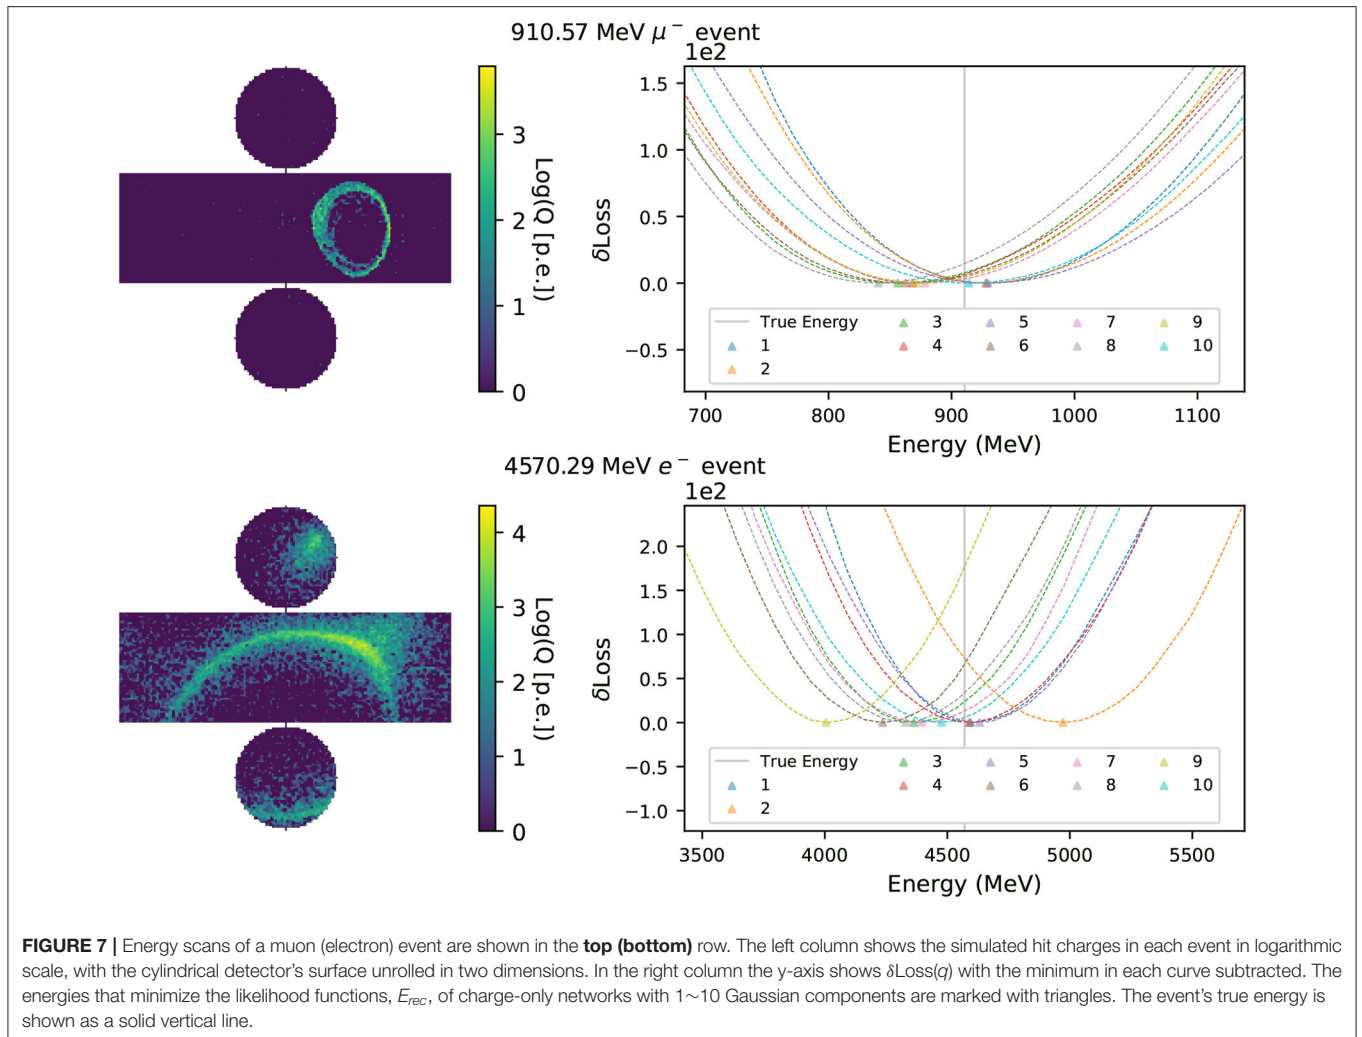


to reproduce sharp features in the data, a shortcoming that is common in generative convolutional neural networks (Durrall et al., 2020). In both cases it is clear the neural network reproduces the general features of the events, including a clear difference between the fuzzy rings predicted for electrons and sharp rings for muons.

To inspect the PDFs for the hit charges and times we have chosen three reference PMTs in the detector, marked in red in **Figure 4**, which are located in the center, edge, and outside of the Cherenkov rings produced by the events generated with fixed kinematics. The PDFs resulting from training the neural network with the charge-only, one-dimensional, loss function using one, three, five and ten components is shown in

**Figure 5**. It is clearly seen that the distributions of simulated hit charges are not Gaussian and therefore the single-Gaussian PDF describes the data very poorly. With three components, the neural network is able to reproduce general features in the data, such as the high-charge tail present in all distributions. With five and particularly with ten components the PDFs describe the distributions in detail, including a slight bi-modality seen at low charge. We have observed that the neural network struggles to reproduce the charge distribution for PMTs on the sharp edge of the muon ring. Given the shape of this distribution is unremarkable compared to all others, we believe this artifact might be due to the known shortcomings of convolutional generative networks to reproduce

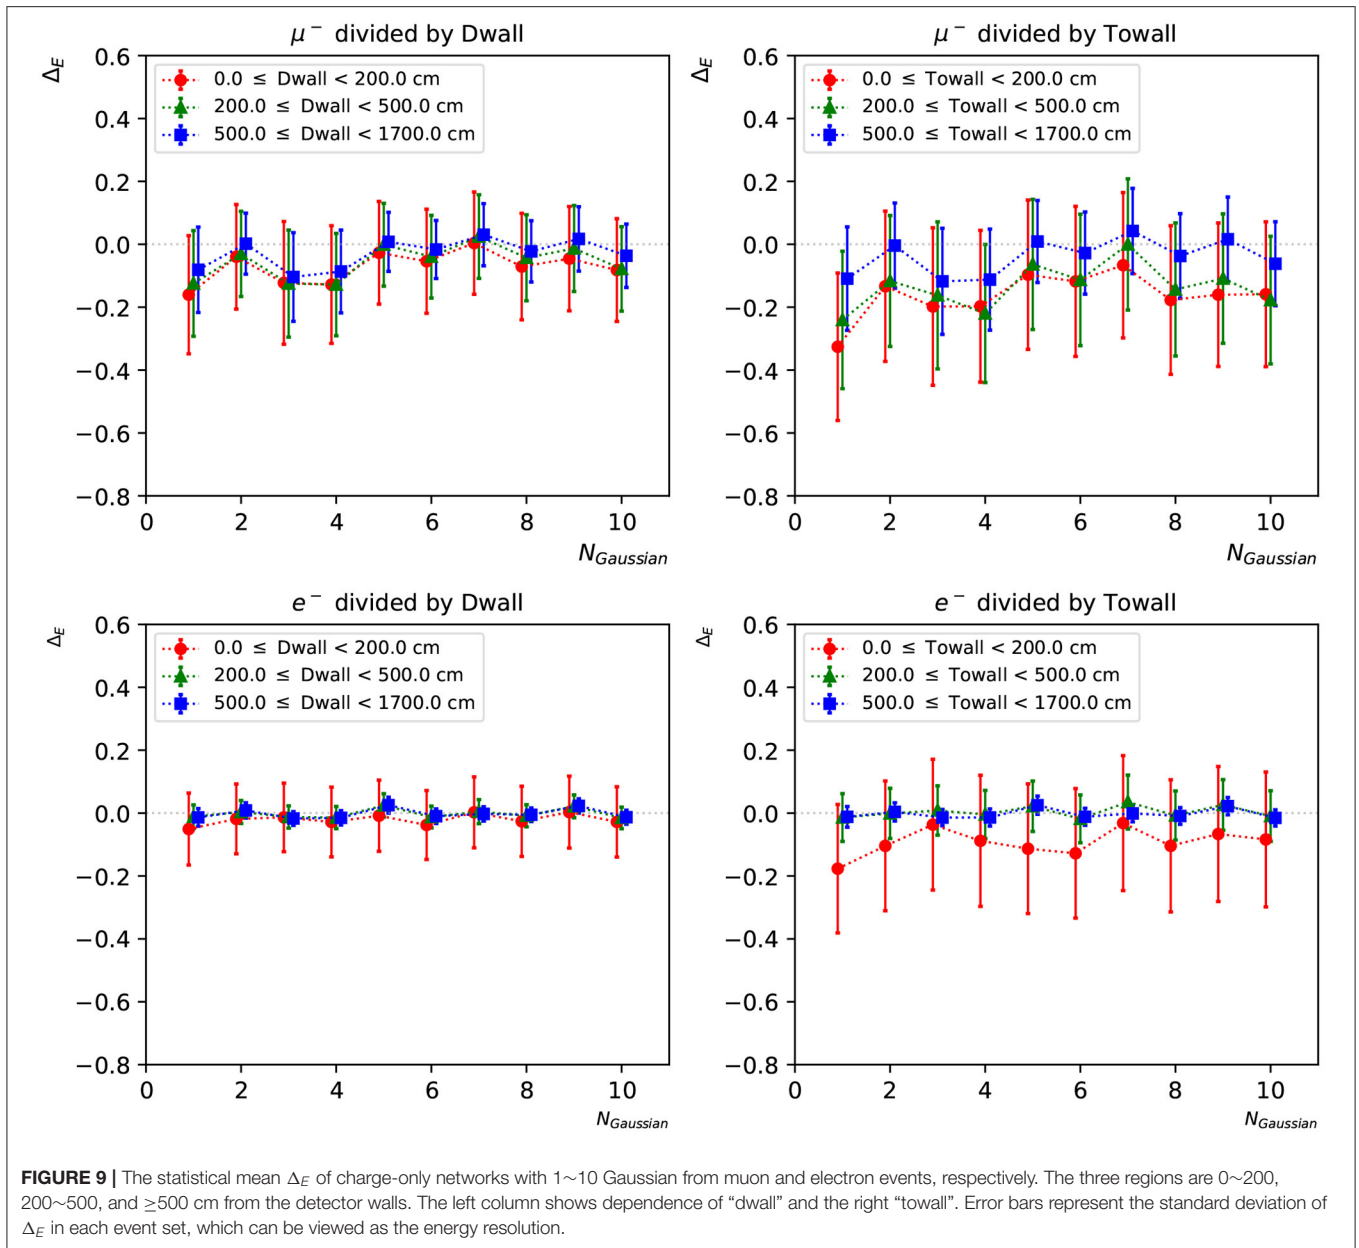




sharp features in the data, as identified in the hit probability prediction above.

The two-dimensional loss function using correlated Gaussian components is inspected similarly for the three chosen PMTs in

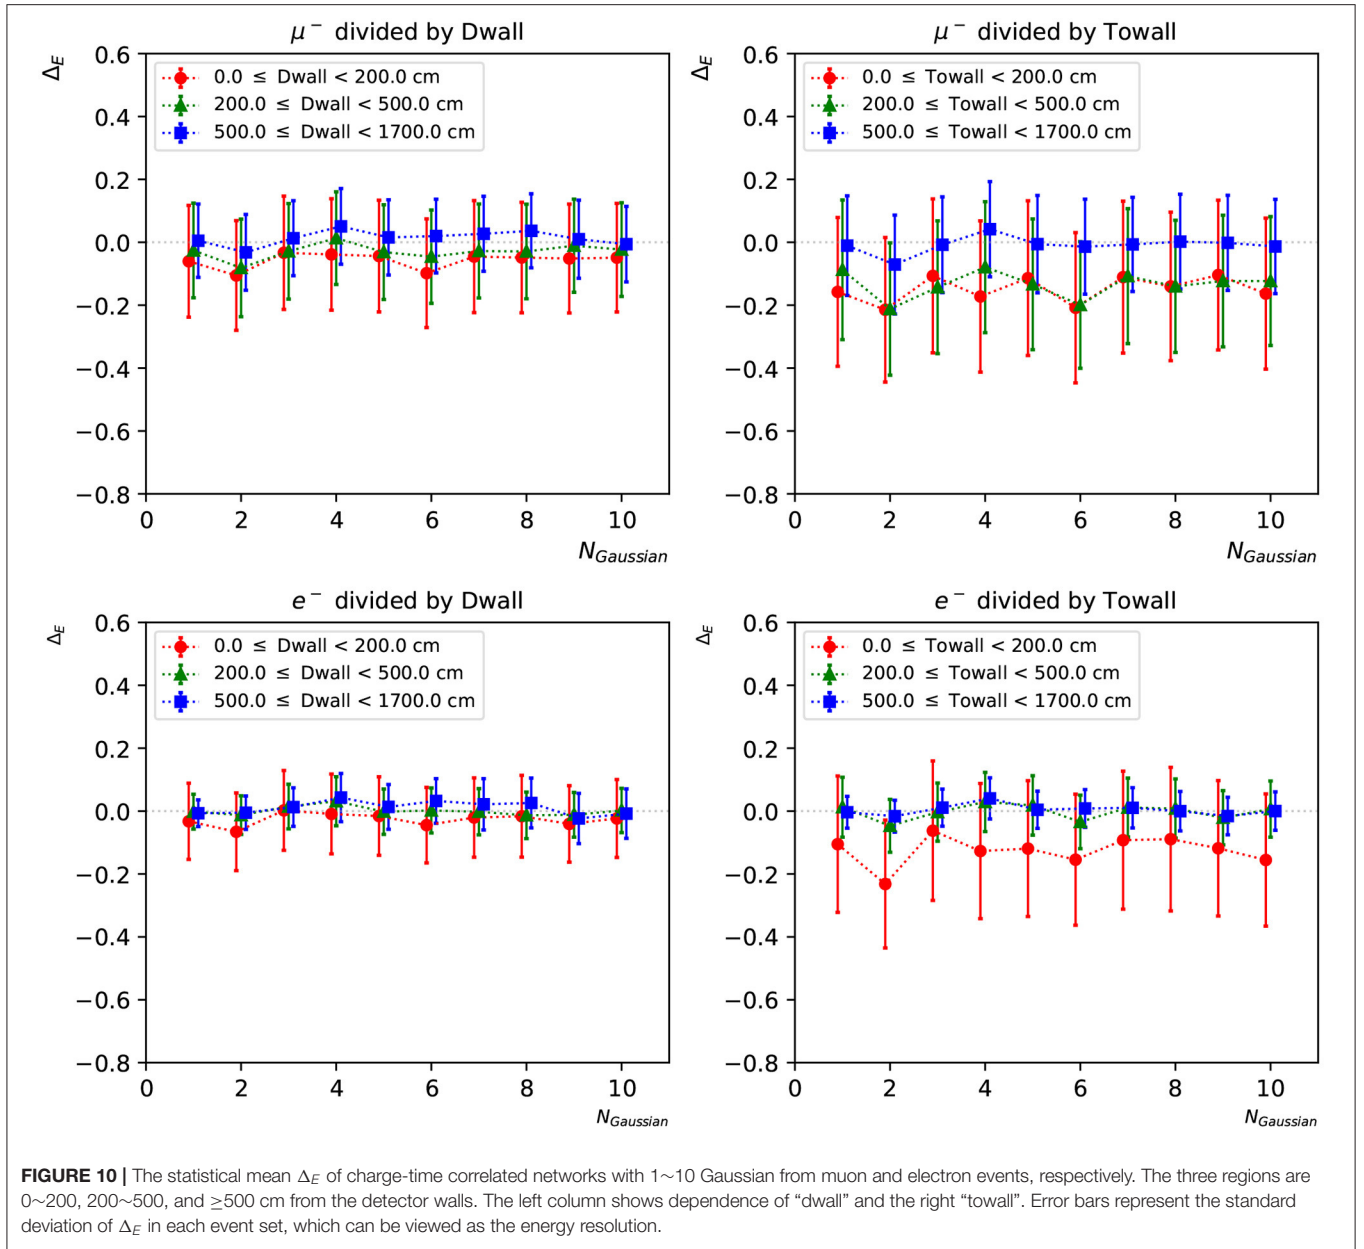
**Figure 6.** The inadequacy of a single two-dimensional Gaussian to describe the data is more evident than in the one-dimensional case. This is expected, given the time distribution is multimodal, with a sharp peak associated to Cherenkov photons



that produce hits without scattering in the water or reflecting off the detector surfaces, and a long tail of scattered and reflected photons, including hints of reflection peaks. The width of the sharp peak is mostly determined by the PMT response, while the distribution of late hits depends strongly on the detector geometry and water properties. A three-component model shows a much better fit, with the network reproducing the multi-modality in the time distribution. The five and ten component models produce increasingly complex shapes with correlations clearly seen in the time and charge dimensions. As in the previous two examples, the neural network's worst performance is for the PMT on the edge of the muon Cherenkov ring, with a bias in the time prediction observed in addition to the relatively poor agreement in the charge prediction.

### 6.2. Likelihood Scans as a Function of Particle Energy

For the method proposed in this work to be viable as an event reconstruction technique, it is important that the likelihood function encoded in the neural network is smooth, so that it can be used in gradient-descent algorithms, and that the minimum of the likelihood surface lies close to the true parameters describing the event. To examine these characteristics of the neural network, we scan the likelihood surface of the randomized data set used for the *in-situ* validation of the neural network, with 100 uniform steps in energy ranging  $-80$  to  $+80\%$  of the event's true energy. The other neural network input parameters remain fixed to their true values. At each scan point, we replace the event's true energy with the shifted one and evaluate the neural network using this hypothesis as the input to calculate the corresponding



loss. A quadruple spline interpolation is applied to the 100 scanned points to find the energy that minimizes the loss, or in other words maximizes the likelihood function. **Figure 7** presents an example of the interpolated energy scan from multi-Gaussian charge-only networks of intermediate energy muon and electron events.

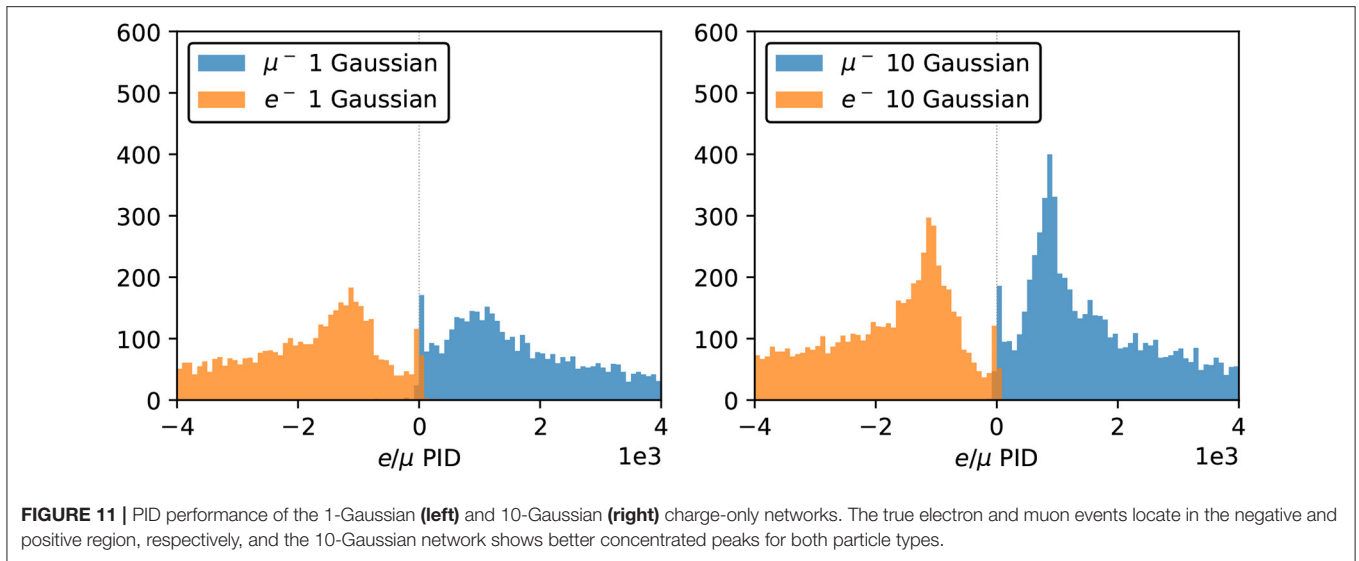
We take the energy which minimizes the loss function to be the estimator for the true particle energy, or the *reconstructed* energy -  $E_{rec}$ , and the fractional energy residual  $\Delta_E$  defined in Equation (9) is used to measure the neural network's energy reconstruction performance.

$$\Delta_E = \frac{E_{rec} - E_{true}}{E_{true}} \quad (9)$$

**Figure 8** shows the distribution of  $\Delta_E$  for sets of 12,000 muon and electron events, using various multi-Gaussian charge-only networks. By the inclusion of the charge-time correlation in each PMT, the small energy biases that exist in the charge-only network reconstruction are significantly improved.

We have applied the likelihood scan to all the three loss functional forms described in Section 3.2 but we show results for the charge-only and charge-time including correlations in this section, because the two networks using PMT time responses have similar reconstruction performance and the correlated version is slightly better.

Under certain circumstances the network will face trouble reconstructing events, especially when not all of the particle energy is deposited inside the detector. These events are more likely to cause reconstruction failure as no valid minimum is



found in the loss function scan. In general the reconstruction failure rate is  $\leq 4\%$  for electron and  $\leq 10\%$  for muon events. Both charge-only and charge-time correlated loss functions have better reconstruction performance for electron events, whereas in the charge-time correlated networks the fraction of unsuccessful scans increases with the number of loss function components.

To further investigate the neural network reconstruction performance, we introduce two parameters: “dwall” is the distance between the particle’s starting position and the nearest detector wall, and “towall” is the distance to the nearest detector wall along the particle’s moving direction. For a particle with small “dwall” and “towall”, it has higher probability to escape the detector and thus not all of the particle energy will be detected. **Figures 9, 10** present the energy reconstruction performance of multi-Gaussian networks in the detector regions defined by the “dwall” and “towall” parameters. Events for which no minimum is found in the loss function scan are excluded. As expected the performance is worse for events with smaller “towall”, with the particle energy being systematically underestimated. The dependence of performance on “towall” is more significant for muons, which tend to travel much longer distances before they drop below the Cherenkov threshold. Due to the longer tails in their  $\Delta_E$  distributions, the muon energy resolution is worse across the three detector regions whereas the electron energy resolution is superior in the central region. For events sufficiently far from the walls, the mean  $E_{rec}$  predicted by any of the 1~10-Gaussian networks is unbiased, with fluctuations smaller than the standard deviation.

### 6.2.1. Particle Identification

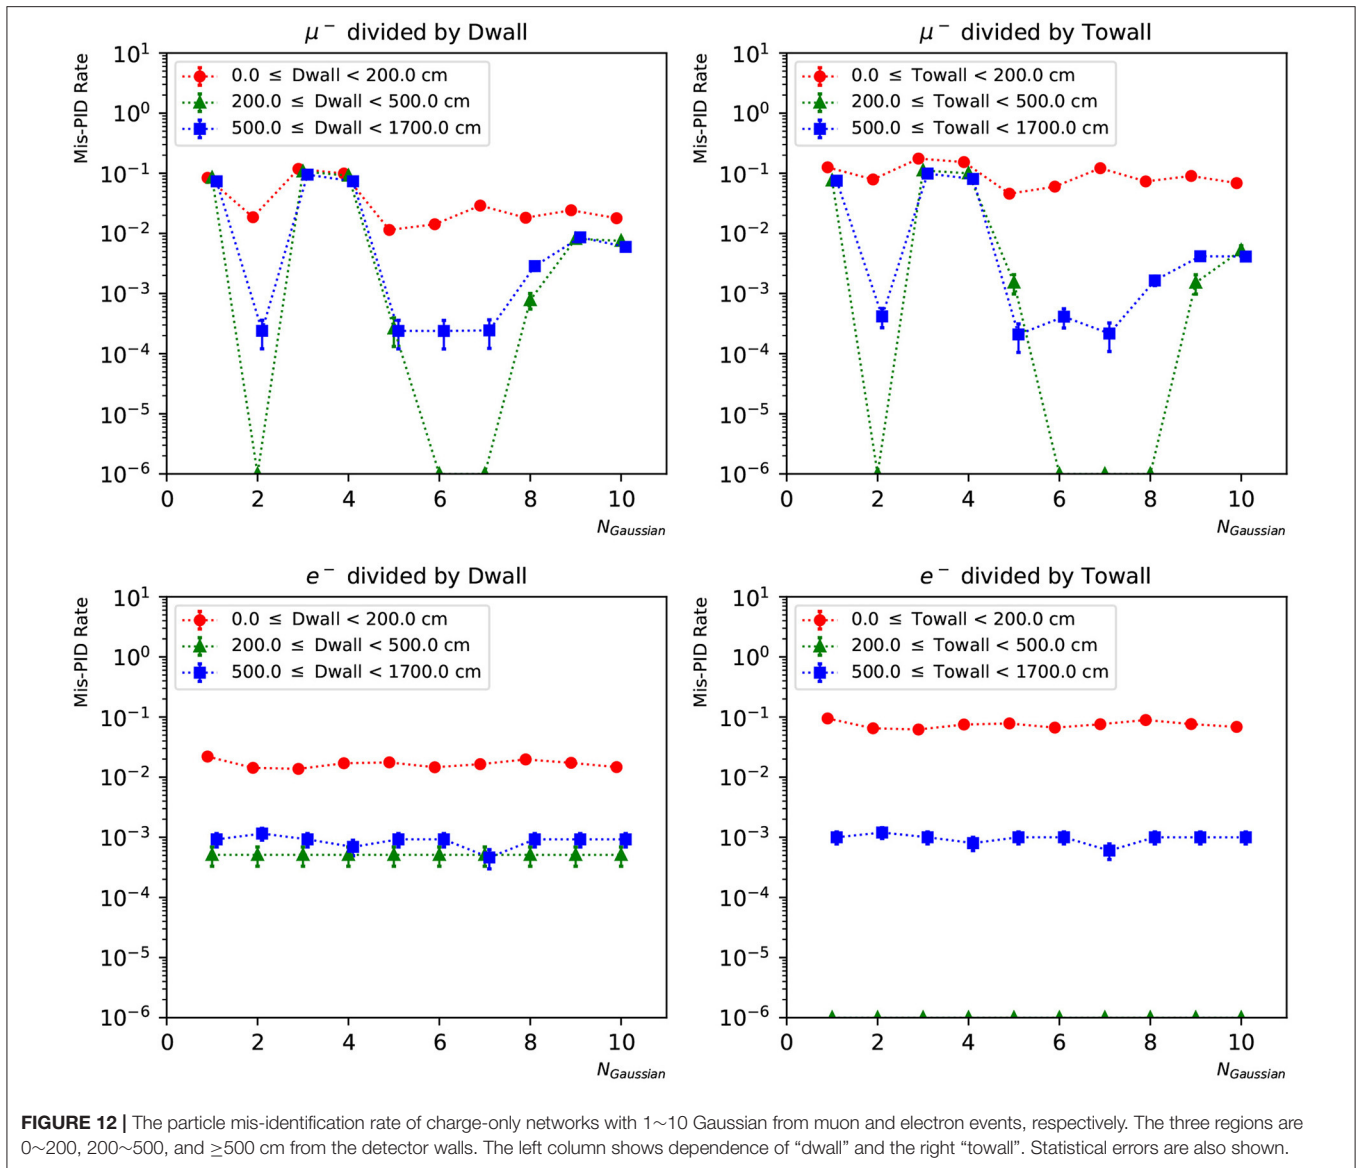
The trained neural networks are able to separate electron and muon events by comparing the loss values of the competing hypotheses. To study the particle identification (PID) performance of the neural network, we use the loss function scans described in Section 6.2 above. For each event, we take the

difference in the loss value at the energy that minimizes the loss,  $E_{rec}$ , while keeping all other input parameters fixed. The resulting variable takes negative values for electron-like events and positive values for muon-like events:

$$e/\mu \text{ PID} = \text{Loss}(q, t, E_{rec})|_{e^-} - \text{Loss}(q, t, E_{rec})|_{\mu^-} \quad (10)$$

**Figure 11** shows the distribution of the  $e/\mu$  PID parameter variable for 12,000 electrons and muons, excluding those events missing a local minimum in the likelihood scan. While the full distribution extends well beyond the plotted range, we focus on the most interesting region, where the two populations cross over, to show the PID performance of our networks. The two particles types are well-separated, with only a small fraction of events crossing the classification boundary at zero. A cluster of events around zero is due to events near the detector walls, which are more difficult to reconstruct as they tend to produce hits in a smaller number of PMTs. We note here that since this PID study is done with all likelihood parameters except energy kept at their true values, the performance shown in this section should be taken as an indicative result. In a realistic event reconstruction setting, all the neural network input parameters will need to be estimated simultaneously by finding the global minimum in the negative log-likelihood surface.

In **Figures 12, 13**, we present the particle mis-identification rate, defined as the wrong-sign  $e/\mu$  PID fraction of each true particle type, with the events broken down in different detector regions. The same conditions as in **Figures 9, 10** are applied. Both the “dwall” and “towall” parameters have strong impact to PID accuracy. In the charge-only networks, muon events sufficiently far from the walls show a noticeable fluctuation of PID accuracy, which is improved by the charge-time correlated networks. Overall both networks exhibit a dependence of particle location for the PID performance. In the same region defined by “dwall” and “towall”, the PID performance of charge-time correlated networks can be improved by including more Gaussian components in the loss function.



**FIGURE 12 |** The particle mis-identification rate of charge-only networks with 1~10 Gaussian from muon and electron events, respectively. The three regions are 0~200, 200~500, and  $\geq 500$  cm from the detector walls. The left column shows dependence of “dwall” and the right “towall”. Statistical errors are also shown.

## 7. DISCUSSION

In this section, we discuss the results described above, as well as observations made during the development of the neural networks presented in this work.

### 7.1. Neural Network Training

While developing the neural network we faced training instability, with the loss function occasionally becoming not-a-number, particularly for more complex functional forms with several components in the Gaussian mixture. We achieved stability by carefully parameterizing the loss functions, as described above, and by pre-scaling our inputs such that they do not significantly exceed unity. Other attempts to stabilize the network training, such as the introduction of batch normalization layers were not as successful.

Especially for simpler versions of the neural networks, we observed that after a rapid decrease of the loss at the start of the training, it would stabilize at an intermediate value for a

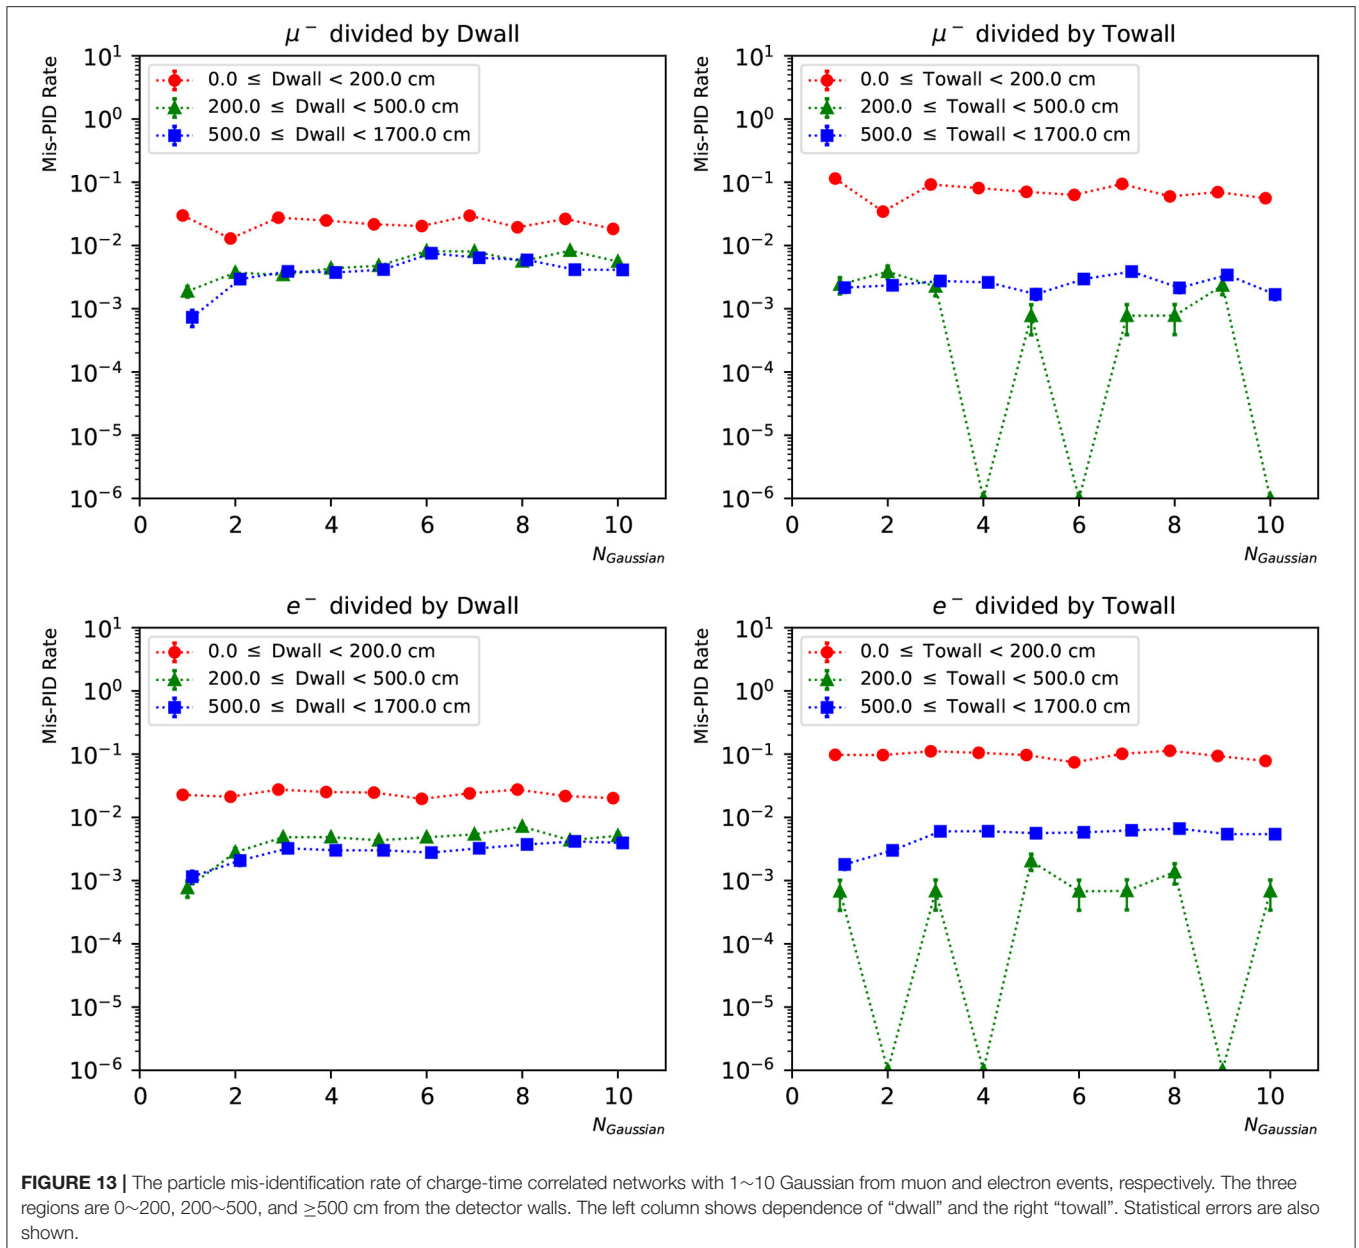
few hundred iterations before decreasing again in a short step to the stable minimum. Inspection of the network output during this period of metastability revealed a clear difference in the generated images before and after the short step, with localized features appearing suddenly after the step, while before the step the network produces a relatively uniform output for the entire image. After the network learns to localize the events, the quality of the generated rings gradually improves with training. This behavior of the loss during training is significantly washed out when using more complex loss functions with several components.

### 7.2. Network Architecture Studies

We tested several modifications to the neural network architecture shown in **Figure 1**, while assessing the performance of the alternative architectures mostly by inspecting the behavior of the loss function during training.

We studied the effects on the network performance of changing the number of fully connected layers and the number





**FIGURE 13** | The particle mis-identification rate of charge-time correlated networks with 1~10 Gaussian from muon and electron events, respectively. The three regions are 0~200, 200~500, and  $\geq 500$  cm from the detector walls. The left column shows dependence of “dwall” and the right “towall”. Statistical errors are also shown.

of nodes in each layer. These studies were done in a version of the network that generated data only for the barrel section of the detector. We changed the number of nodes in each fully connected layer from three to 1,000 and found that 30 nodes were sufficient, and adding more nodes did not improve the results. We found that at least two fully connected layers (excluding the input layer) were necessary to generate events on the barrel.

In addition, we studied the effect of the changing the depth of the network’s convolutional layers using the single-Gaussian charge-only loss function. Depths of 32, 64, 128, and 256 were tested, where 64 is the default number. We found that increasing the depth makes the network settle faster, while achieving a more or less similar loss value after training. An example of such a study is shown in **Figure 14**.

We also tested using three different activation functions throughout the network: *ReLU*, *tanh*, and *LeakyReLU* where the  $f(x) = 0$  portion of *ReLU* is replaced by  $f(x) = mx$ . The *LeakyReLU* activation function was examined with two different negative  $x$  slopes,  $m = 0.1$  and  $m = 0.5$ . The *tanh* and *LeakyReLU* performed equally well or worse compared to *ReLU*.

### 7.3. Network Reconstruction Performance

With the results shown in Sections 6.1 and 6.2, we notice that the multi-Gaussian PDF can improve the reconstruction of individual PMT responses, but has limited impact to the particle event reconstruction, possibly due to insufficient training. With the inclusion of PMT time responses, the charge-time correlated networks show larger reconstruction uncertainty than the charge-only ones,

especially for electrons. We have not determined what causes this and we will investigate this effect further.

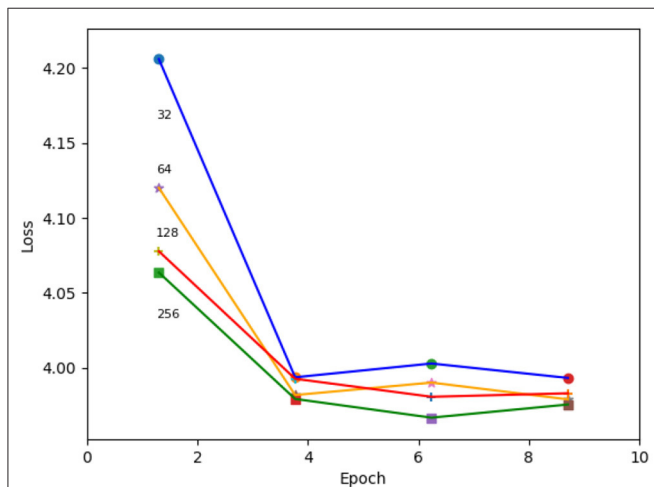
Moreover, event reconstruction performance strongly depends on the event’s location, as a near-wall one is less likely to deposit all of its energy in the detector, especially in the case of muons. In addition, the Cherenkov ring is sampled by fewer PMTs for events with small “towall”. **Figure 15** shows the distribution of  $\Delta_E$  in the plane of  $E_{rec}$  and “towall” for muon and electron respectively. In the region where a particle’s energy is fully captured by the detector, e.g., low  $E_{rec}$  and high “towall”, the network shows excellent energy reconstruction accuracy. A linear relation is noticeable between muon’s  $E_{rec}$

and “towall”, which is consistent with the expectation: muon track lengths are approximately proportional to their energy and once a muon track escapes the detector it is no longer possible to accurately measure its energy. Water Cherenkov experiments typically include external detector elements to identify escaping events that cannot be reliably reconstructed. Electron showers, on the other hand, have a much more limited extent, and therefore their energy is accurately reconstructed even at high energies and relatively small values of “towall”. Another interesting observation in these results is the smooth appearance of well-reconstructed muon events at 100~200 MeV as shown in **Figure 15**. These energies correspond roughly to the Cherenkov threshold and the smooth transition indicates the neural network is capable of identifying muon rings close to the threshold, albeit with a limited efficiency. Detecting particles close to their threshold is important particularly when considering events with multiple particle topologies, where very faint rings can coexist with much brighter rings, making their identification challenging.

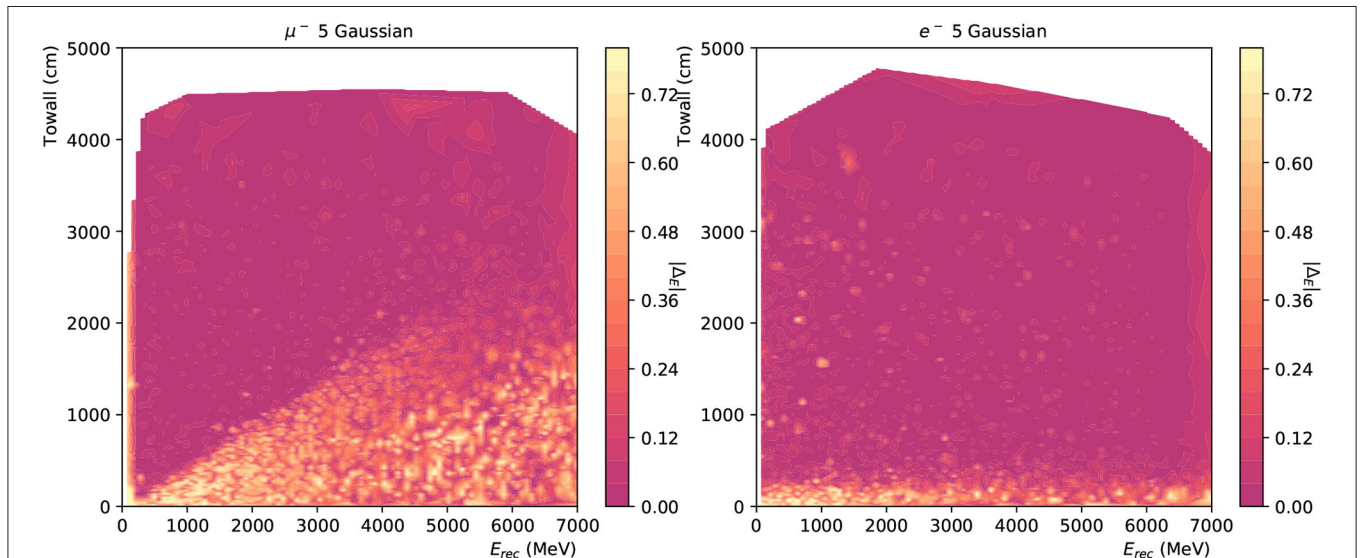
While we observe some dependence of the neural network performance on the number of Gaussian components used for the reconstruction, identifying the optimal number of components will require further study. In particular, it will be important to measure the neural network performance in a more realistic event reconstruction scenario, such as finding the likelihood maxima using gradient descent, and considering also the event reconstruction runtime.

## 8. CONCLUSION

We have developed a deep convolutional neural network that generates the likelihood function for the complete set of observables in a cylindrical water Cherenkov detector. We



**FIGURE 14** | Loss as a function of training epoch for varying depths of the neural network’s convolutional layers. The depth of the convolutional layers used for each training is indicated at the start of the respective line.



**FIGURE 15** | The energy reconstruction performance of 12,000 muon (left) and electron (right) events plotted against the  $E_{rec}$  and towall for charge-only network with 5 Gaussian. The saturation level at 0.8 is from the energy range for likelihood scans.

presented results using three likelihood function approximations using mixtures of Gaussian functions in one-dimension, taking into account only the hit charges, and in two-dimensions, taking into account both the hit charges and times. For the bivariate mixture model we explored two parameterizations of the Gaussian distributions: one which does not allow for correlations between charge and time in the individual components of the mixture, and one which includes this additional degree of freedom.

We have demonstrated that our neural networks are capable of accurately reproducing the features observed in the data and we found our results to be robust with respect to changes in the network architecture and loss function parameterizations, though longer training of our neural network will be necessary to confirm some of our conclusions.

We have studied the potential of these neural networks to be used in maximum likelihood reconstruction by scanning the loss as a function of the particle energy for a set of muon and electron events, and we demonstrated the loss is minimized close to the true particle energy. Taking the difference of between the minimum loss under the electron hypothesis and the minimum loss under the muon hypothesis to form a likelihood ratio test, the neural network has shown very promising particle identification performance.

The work presented here is an initial milestone in the path to achieve precise water Cherenkov event reconstruction using ML-based maximum likelihood estimation. Our future work will focus on the application of gradient descent algorithms to the neural networks presented here, and further improving their performance based on event reconstruction metrics. We also plan to continue exploring alternative loss function parameterizations, such as mixtures of log-normal and other simple functions. We foresee the next major milestone in the project to be the extension of the neural network to multiple-particle topologies, which we expect to achieve by combining several single-particle likelihoods, possibly using ML techniques.

## REFERENCES

- Abe, K., et al. (2020). Constraint on the matter-antimatter symmetry-violating phase in neutrino oscillations. *Nature* 580, 339–344. doi: 10.1038/s41586-020-2177-0
- Abe, K., Abe, K., Aihara, H., Aimi, A., Akutsu, R., Andreopoulos, C., et al. (2018). Hyper-Kamiokande design report. *arXiv [Preprint] arXiv:1805.04163*. doi: 10.48550/arXiv.1805.04163
- Abi, B., et al. (2020). Neutrino interaction classification with a convolutional neural network in the DUNE far detector. *Phys. Rev. D* 102, 092003. doi: 10.1103/PhysRevD.102.092003
- Abratenko, P., et al. (2021). Convolutional neural network for multiple particle identification in the MicroBooNE liquid argon time projection chamber. *Phys. Rev. D* 103, 092003. doi: 10.1103/PhysRevD.103.092003
- Agostinelli, S., et al. (2003). GEANT4-a simulation toolkit. *Nuclear Instrum. Methods* 506, 250–303. doi: 10.1016/S0168-9002(03)01368-8
- Alanazi, Y., Sato, N., Ambrozewicz, P., Blin, A. N. H., Melnitchouk, W., Battaglieri, M., et al. (2021). “A survey of machine learning-based physics event generation,” in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence* (Montreal, QC). doi: 10.24963/ijcai.2021/588

The approach presented here is not only promising in terms of potential improvements in reconstruction performance, but it is also easily applicable to any cylindrical water Cherenkov experiment, with extension to other geometries possible as long as they can be reasonably projected onto a set of two-dimensional images.

## DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

## AUTHOR CONTRIBUTIONS

The study was initiated by CV and led by CV, MW, and CY. The data was prepared by MJ and CV. MJ, KK, LM, AP, CV, JX, and CY contributed to the design of the neural network architecture. MJ, CV, JX, and KY contributed to the reconstruction performance studies. All authors contributed to manuscript revision, read, and approved the submitted version.

## ACKNOWLEDGMENTS

We acknowledge the support of the Water Cherenkov Machine Learning (WatChMaL) group<sup>4</sup>, particularly in the early stages of the project. We would like to thank Stony Brook Research Computing and Cyberinfrastructure, and the Institute for Advanced Computational Science at Stony Brook University for access to the high-performance SeaWulf computing system, which was made possible by a \$1.4M National Science Foundation grant (#1531492). This research was enabled in part by support provided by WestGrid (<https://www.westgrid.ca/support/>) and Compute Canada ([www.computecanada.ca](http://www.computecanada.ca)) with its high-performance heterogeneous cluster Cedar.

<sup>4</sup><https://www.watchmal.org/>

- Aleku, A., et al. (2021). Updated physics performance of the ESSnuSB experiment: ESSnuSB collaboration. *Eur. Phys. J.* 81, 1130. doi: 10.1140/epjc/s10052-021-09845-8
- Alonso-Monsalve, S., Douqa, D., Jesús-Valls, C., Lux, T., Pina-Otey, S., Sánchez, F., et al. (2021). Graph neural network for 3D classification of ambiguities and optical crosstalk in scintillator-based neutrino detectors. *Phys. Rev. D* 103, 032005. doi: 10.1103/PhysRevD.103.032005
- Alvarez-Ruso, L., et al. (2018). NuSTEC White Paper: status and challenges of neutrino-nucleus scattering. *Prog. Part. Nucl. Phys.* 100, 1–68. doi: 10.1016/j.pnpnp.2018.01.006
- Askin, M., Bagdasarian, Z., Barros, N., Beier, E., Blucher, E., Bonventre, R., et al. (2020). Theia: an advanced optical neutrino detector. *Eur. Phys. J.* 80, 416. doi: 10.1140/epjc/s10052-020-7977-8
- Aurisano, A., Radovic, A., Rocco, D., Himmel, A., Messier, M. D., Niner, E., et al. (2016). A convolutional neural network neutrino event classifier. *J. Inst.* 11, P09001. doi: 10.1088/1748-0221/11/09/P09001
- Becker-Szendy, R., Bionta, R., Bratton, C., Casper, D., Claus, R., Cortez, B., et al. (1993). A large water Cherenkov detector for nucleon decay and neutrino interactions. *Nuclear Instrum. Methods* 324, 363–382. doi: 10.1016/0168-9002(93)90998-W

- Brun, R., and Rademakers, F. (1997). ROOT: an object oriented data analysis framework. *Nuclear Instrum. Methods* 389, 81–86. doi: 10.1016/S0168-9002(97)00048-X
- Butter, A., and Plehn, T. (2020). Generative networks for LHC events. *arXiv [Preprint] arXiv:2008.08558*. doi: 10.48550/arXiv.2008.08558
- Cover, T., and Thomas, J. (1991). *Elements of Information Theory*. New York, NY: Wiley. doi: 10.1002/0471200611
- Dosovitskiy, A., Springenberg, J. T., and Brox, T. (2014). Learning to generate chairs with convolutional neural networks. *arXiv [Preprint] arXiv:1411.5928*. doi: 10.1109/CVPR.2015.7298761
- Durall, R., Keuper, M., and Keuper, J. (2020). “Watch your up-convolution: CNN based generative deep neural networks are failing to reproduce spectral distributions,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (IEEE)*, 7890–7899. Available online at: [https://openaccess.thecvf.com/content\\_CVPR\\_2020/papers/Durall\\_Watch\\_Your\\_Up\\_Convolution\\_CNN\\_Based\\_Generative\\_Deep\\_Neural\\_Networks\\_Are\\_CVPR\\_2020\\_paper.pdf](https://openaccess.thecvf.com/content_CVPR_2020/papers/Durall_Watch_Your_Up_Convolution_CNN_Based_Generative_Deep_Neural_Networks_Are_CVPR_2020_paper.pdf)
- Fukuda, S., Fukuda, Y., Hayakawa, T., Ichihara, E., Ishitsuka, M., Itow, Y., et al. (2003). The Super-Kamiokande detector. *Nuclear. Instrum. Methods* 501, 418–462. doi: 10.1016/S0168-9002(03)00425-X
- Fukuda, Y., Hayakawa, T., Ichihara, E., Inoue, K., Ishihara, K., Ishino, H., et al. (1998a). Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.* 81, 1562–1567.
- Fukuda, Y., Hayakawa, T., Ichihara, E., Inoue, K., Ishihara, K., Ishino, H., et al. (1998b). Measurements of the solar neutrino flux from Super-Kamiokande's first 300 days. *Phys. Rev. Lett.* 81, 1158–1162.
- Hirata, K., Kajita, T., Koshiba, M., Nakahata, M., Oyama, Y., Sato, N., et al. (1987). Observation of a neutrino burst from the supernova SN 1987a. *Phys. Rev. Lett.* 58, 1490–1493. doi: 10.1103/PhysRevLett.58.1490
- James, F. (1994). *MINUIT Function Minimization and Error Analysis: Reference Manual Version 94.1*. Geneva.
- Jiang, M., et al. (2019). Atmospheric neutrino oscillation analysis with improved event reconstruction in Super-Kamiokande IV. *Prog. Theoret. Exp. Phys.* 2019, 053F01. doi: 10.1093/ptep/ptz015
- Kingma, D. P., and Ba, J. (2017). Adam: a method for stochastic optimization. *arXiv [Preprint] arXiv:1412.6980*. doi: 10.48550/arXiv.1412.6980
- Neyman, J., and Pearson, E. S. (1933). On the problem of the most efficient tests of statistical hypotheses. *Philos. Trans. R. Soc. Lond. A* 231, 289–337. doi: 10.1098/rsta.1933.0009
- Oyama, Y., Hirata, K., Kajita, T., Kifune, T., Kihara, K., Nakahata, M., et al. (1989). Experimental study of upward-going muons in Kamiokande. *Phys. Rev. D* 39, 1481–1491. doi: 10.1103/PhysRevD.39.1481
- Paganini, M., de Oliveira, L., and Nachman, B. (2018). CaloGAN: Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks. *Phys. Rev. D* 97, 014021. doi: 10.1103/PhysRevD.97.014021
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). “PyTorch: an imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems*, Vol. 32, eds H. Wallach, H. Larochelle, A. Beygelzimer, F. D’Alche-Buc, E. Fox, and R. Garnett (Vancouver, BC: Curran Associates, Inc.), 8024–8035.
- Perdue, G. N., et al. (2018). Reducing model bias in a deep learning classifier using domain adversarial neural networks in the MINERvA experiment. *J. Inst.* 13, P11020. doi: 10.1088/1748-0221/13/11/P11020
- Radovic, A., Williams, M., Rousseau, D., Kagan, M., Bonacorsi, D., Himmel, A., et al. (2018). Machine learning at the energy and intensity frontiers of particle physics. *Nature* 560, 41–48. doi: 10.1038/s41586-018-0361-2
- Takenaka, A., Abe, K., Bronner, C., Hayato, Y., Ikeda, M., Imaizumi, S., et al. (2020). Search for proton decay via  $p \rightarrow e^+\pi^0$  and  $p \rightarrow \mu^+\pi^0$  with an enlarged fiducial volume in Super-Kamiokande I-IV. *Phys. Rev. D* 102, 112011. doi: 10.1103/PhysRevD.102.112011
- Vallecorsa, S., Carminati, F., and Khattak, G. (2019). 3D convolutional GAN for fast simulation. *EPJ Web Conf.* 214, 02010. doi: 10.1051/epjconf/201921402010
- Williams, P. M. (1996). Using neural networks to model conditional multivariate densities. *Neural Comput.* 8, 843–854. doi: 10.1162/neco.1996.8.4.843
- Xia, J., Nishimura, Y., Bronner, C., and Mochizuki, T. (2019). The calibration and evaluation of 50cm box and line photomultiplier tubes designed for Hyper-Kamiokande. *JPS Conf. Proc.* 27, 012002. doi: 10.7566/JPSCP.27.012002
- Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.
- Publisher’s Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.
- Copyright © 2022 Jia, Kumar, Mackey, Putra, Vilela, Wilking, Xia, Yanagisawa and Yang. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.