

Research Article

Deep Learning the Effects of Photon Sensors on the Event Reconstruction Performance in an Antineutrino Detector

Chang-Wei Loh, Zhi-Qiang Qian, Rui Zhang, You-Hang Liu, De-Wen Cao, Wei Wang, Hai-Bo Yang, and Ming Qi 

Physics Department, Nanjing University, 22 Hankou Road, Nanjing, Jiangsu, China

Correspondence should be addressed to Ming Qi; mqi@cern.ch

Received 3 January 2018; Revised 3 April 2018; Accepted 15 April 2018; Published 9 July 2018

Academic Editor: Hiroyasu Ejiri

Copyright © 2018 Chang-Wei Loh et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The publication of this article was funded by SCOAP³.

We provide a fast approach incorporating the usage of deep learning for studying the effects of the number of photon sensors in an antineutrino detector on the event reconstruction performance therein. This work is a first attempt to harness the power of deep learning for detector designing and upgrade planning. Using the Daya Bay detector as a case study and the vertex reconstruction performance as the objective for the deep neural network, we find that the photomultiplier tubes (PMTs) at Daya Bay have different relative importance to the vertex reconstruction. More importantly, the vertex position resolutions for the Daya Bay detector follow approximately a multiexponential relationship with respect to the number of PMTs and, hence, the coverage. This could also assist in deciding on the merits of installing additional PMTs for future detector plans. The approach could easily be used with other objectives in place of vertex reconstruction.

1. Introduction

The choice of photon sensors such as photomultiplier tubes (PMTs), be it their expected sizes, locations, and the total number of sensors in antineutrino detectors, including Daya Bay [1], Double Chooz [2], RENO [3], and JUNO [4], are of interest as these sensors are the information gatherers through which we can identify antineutrino interaction events. This work is an attempt in using machine learning, in particular deep learning [5–7] as a way to understand how the number of PMTs in the detector influences the event reconstruction performance and extract lessons to be learned therefrom for areas such as detector designing and upgrade planning. To the best of our knowledge, this work is the first study on the efficacy of deep learning in detector designing and planning. For this work, we ask the following: suppose we are given N possible number of locations for the installation of k number of PMTs ($k \leq N$), where should the PMTs be installed such that the event interaction vertex reconstruction is optimal or near-optimal given only these k PMTs in the detector and N possible locations? Of course, N and k could be infinite, but this is technically impossible as it would not

meet the budget of a detector construction. In this work, we use deep learning on a model of the Daya Bay detector as a case study to understand the impact of PMTs on event position vertex reconstructions in a detector. The vertex is useful for studies on signal-background discriminations and the correction to the position-dependent energy response in the detector. The reconstruction of the vertices has been studied in-depth in the Daya Bay experiment using non-machine learning methods. As such, this allows us to cross-check our vertex reconstruction with deep learning with other methods in the Daya Bay before studying its potential in detector designing. Moreover, the vertex reconstruction performance is chosen as the objective since it is relatively simple for deep learning to handle for a clear understanding of the approach without involving too much experimental details. Nonetheless, experimentalists can easily substitute the vertex reconstruction performance with other objectives of interest. Beyond antineutrino detectors, sensor placements have been studied in areas ranging from water network distributions [8] to fault detections [9].

The Daya Bay antineutrino detectors are liquid scintillator detectors with a physics program focusing on the

precision measurement of the neutrino mixing angle θ_{13} with reactor antineutrinos. Each Daya Bay detector consists of three concentric cylindrical tanks: an inner acrylic vessel (IAV) containing gadolinium- (Gd-) doped liquid scintillator, an outer acrylic vessel (OAV) containing undoped liquid scintillator which surrounds the IAV, and a stainless steel vessel (SSV) which surrounds the IAV and OAV. With this design, the detectors could detect the interaction of the antineutrinos and the scintillator via inverse beta decay (IBD) reactions:

$$\bar{\nu}_e + p \longrightarrow e^+ + n. \quad (1)$$

The emitted positron then undergoes ionization processes in the liquid scintillator before annihilating with an electron producing a prompt signal with an energy deposition in the range of 1 - 8 MeV. The deposited energy is converted to scintillation photons which are then collected by the PMTs. As the positron displacement prior to the annihilation is negligible, the interaction vertex of the prompt signal can be assumed to be the antineutrino IBD interaction vertex. However, the neutron thermalizes and diffuses before being captured on either a proton or Gd with a mean capture time of $\sim 30 \mu\text{s}$ in the Gd-doped liquid scintillator and $\sim 200 \mu\text{s}$ in the undoped liquid scintillator, giving rise to a delayed signal. A total of 192 Hamamatsu R5912 8-inch PMTs [10], arranged in a layout with 8 rows and 24 columns, are installed on the vertical wall of the SSV pointing inward towards the OAV and IAV forming a total of 6% photodetector coverage. Located above and below the OAV are reflective panels that serve to redirect scintillation light towards the PMTs thereby increasing the photon collection efficiency to 12% effectively.

As aforementioned, we used deep learning to perform the IBD vertex reconstruction in order to study the effects of PMTs on an event reconstruction. Deep learning is a class of machine learning, which is especially adept at leveraging large datasets to compute human-comprehensible quantities by learning the various degrees of correlations within. Notably, it can, on its own, learn to discover functional relationships from the data without *a priori* given, effectively forming a mapping from the inputs to a quantity of interest. In other words, deep learning seeks to model the quantity of interest y using a vector of inputs x with $DL(x, p) = y$, where p are parameters of the deep network; their numerical values were found by minimizing the error between the predicted y' and y . Deep learning machine architectures, commonly known as deep neural networks (DNN), are based on artificial neural networks [11] but deeper in terms of the number of hidden layers and are more flexible in terms of how each neuron is connected to other neurons.

The ubiquity of deep learning and its significant success over traditional methods across disparate fields [12–14] in discovering patterns is surprising. However, this may well be due, in part, to that our universe operates on simple physical properties [15]. In high-energy physics, deep learning has demonstrated its prospective use in jets [16, 17], as part of the signal-background discrimination toolkit in the search for beyond the Standard Model particles [18] and Higgs bosons [19] and in neutrino physics experiments [20–24].

2. Recursive Search

As mentioned in Section 1, we wish to search for the k most important locations corresponding to k PMTs installed therein from the N total possible locations in the detector in determining the vertex position V of events collected from the detector. N is a free parameter which could be chosen during the detector design and simulation stage. Denoting the set containing the k number of most important locations as the set S_k^* , this implies that we should find the set S_k^* such that the vertex reconstruction error is minimized. However, finding such k locations simultaneously is a task confounded by a computation that grows exponentially with k . Alternatively, we could search for an approximation to S_k^* by recursively finding the important PMT location one at a time, which can be achieved using deep learning. Since searching for the most important location is equivalent to searching for the most important PMT at that particular location, the phrase “ k -th important PMT” will be used in this work as a shorthand for “ k -th important PMT location”.

Let the true position of the IBD prompt events be $V_{true} = \{x_{true}, y_{true}, z_{true}\}$; the predicted position using DNN as $V_{pred} = \{x_{pred}, y_{pred}, z_{pred}\}$, then in a recursive search, the k -th important PMT, PMT_k^* , will be the one that maximizes the improvement in the resolution σ of the residual distribution given that the $(k - 1)$ other PMTs have already been found through the recursive search, i.e., $\{PMT_{k-1}^*, PMT_{k-2}^*, \dots, PMT_1^*\}$, and where the residual is $V_{pred} - V_{true}$. Namely,

$$PMT_k^* = \arg \min_{P_k \in N \setminus S_{k-1}^{recu}} \sigma(PMT_k, S_{k-1}^{recu}), \quad (2)$$

where $S_{k-1}^{recu} = \{PMT_{k-1}^*, \dots, PMT_1^*\}$ and N is the set containing all the PMTs. Using (2), PMTs could be progressively added into a larger and larger subset S defining the best set found by the algorithm. Alternatively, one could perform a backward elimination: starting from the set with all PMTs and progressively eliminating the most “unimportant” PMT. At the conclusion of this recursive search, we obtain a curve of the event reconstruction resolution versus the number of PMTs used for the reconstruction thereof.

3. Deep Neural Network

In our approach utilizing DNNs, we used a Monte Carlo dataset comprising 2 million IBD prompt events obtained from a Daya Bay detector model which were randomly partitioned into a training set (1.4 million), a validation set (0.3 million), and a test set (0.3 million). The validation set is used for the early stopping of the DNN training to prevent overfitting or underfitting of the data [25]. The parameter N as defined in Section 2 would be 192 corresponding to the 192 PMT locations in the Daya Bay detector model. The charge information of the 192 PMTs is fed into the DNN as its inputs, and the output is the predicted vertex location V_{pred} . To train the DNN, we used the mean square error (*MSE*) loss function

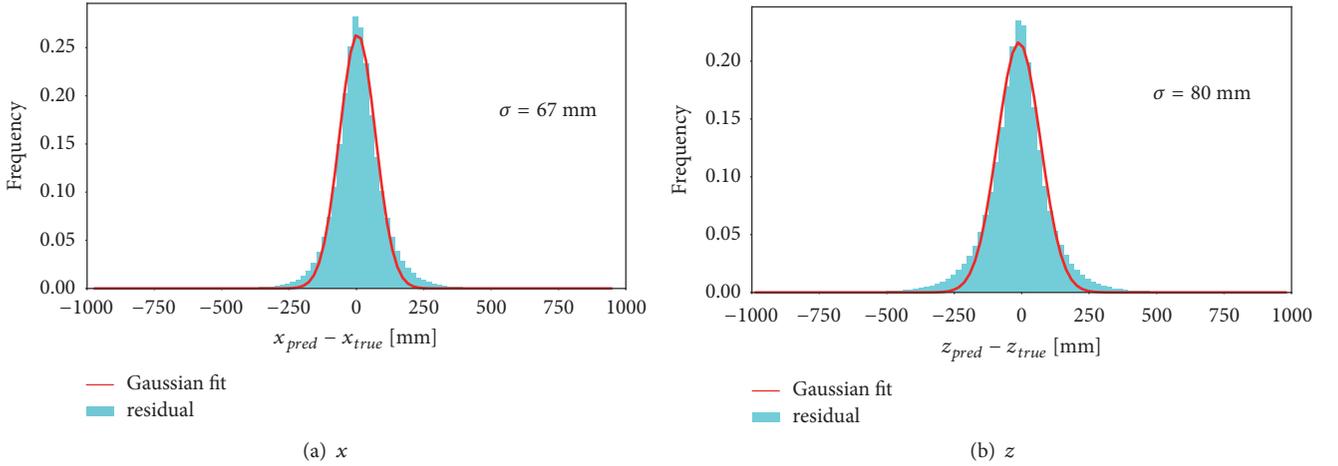


FIGURE 1: Residual distributions for x and z using all 192 PMT charge information.

to measure the error between the predicted, and the truth vertex positions:

$$MSE = \frac{1}{T} \sum_{j=1}^T \left(v_{ij_{pred}} - v_{ij_{true}} \right)^2, \quad (3)$$

where T is the number of events, $v_{ij_{pred}}$ and $v_{ij_{true}}$ are the predicted and truth values for the i -th coordinate of the j -th event vertex, respectively ($i = x, y, z$). The MSE was minimized to obtain the optimal DNN parameters. The minimization is typically done with a gradient descent method [26] involving the gradient of the loss function with respect to the DNN parameters, including the weights of each neuron; i.e., at each training iteration, the parameters w are updated via

$$w_i = w_i - \gamma \cdot \frac{\partial MSE}{\partial w_i}, \quad (4)$$

where γ is the learning rate determined by the user that controls the step length in the negative gradient direction during the training stage. When the MSE reaches minima, $\partial MSE / \partial w_i = 0$. At this point, the DNN has found the needed parameter values to best reconstruct the vertex position. To train the DNN, γ starts with a value of 0.001 and is progressively multiplied a factor of 0.5 whenever the value of the loss function metric stops improving. In this manner, the DNN training will descent quickly in the direction of the minima in the early stage; with a smaller learning rate at a later stage, the training will not overshoot the minima but will descent steadily towards it. An early stopping is made during the training, whereby the training is terminated when no further improvements could be observed from the loss function value after a predetermined number of training rounds, in this case ten. Without such early stopping, the loss function value can rise again indicating that an overfitting has occurred.

The efficacy of deep learning to predict the position of the IBD prompt events can be demonstrated by the residual distributions shown in Figure 1 where the charge

information of the 192 PMTs is fed into a DNN as its inputs. The DNN used here to obtain V_{pred} consists of multiple fully connected layers with ReLU [27] hidden neurons. The optimal number of layers and neurons were obtained using a tree-structured Parzen estimator [28]. The resulting network comprises three hidden layers containing 180, 148, and 148 neurons, respectively. The resolutions as obtained from the Gaussian fit to the residual distributions are 67 mm and 80 mm for $(x_{pred} - x_{true})$ and $(z_{pred} - z_{true})$, respectively.

A straightforward and brute force use of (2) in a recursive search using a DNN to identify PMT_k^* would be to check over all the remaining PMTs not in the optimal set S_{k-1}^* and separately construct the residual distributions, picking the one giving the best resolution for a particular coordinate in V_{true} . For this brute force search, we used a DNN architecture similar to the aforementioned DNN. The input layer will contain neurons with charge information from the already-chosen PMTs, i.e., those in S_{k-1}^{recu} , plus a candidate PMT, i.e., PMT_k . The computation time for such a search grows quadratically with the total number of PMTs in the detector. Such a brute force search is clearly not scalable. Hence, in this work, we have also used a fast approach to approximate the brute force search but which mitigates the nonscalability of the latter.

This fast approach integrates a DNN component from the autoencoder architecture [29]: a bottleneck layer with a single neuron, as shown in Figure 2. In this bottleneck DNN architecture, the remaining candidate PMTs not in S_{k-1}^{recu} are forced to connect to the bottleneck neuron before being given to the fully connected layers as inputs, effectively demanding the DNN to search for the best weights associated with each of these PMTs. At the bottleneck region, the DNN computes the sum $\sum w_i PMT_i$, where i runs over the candidate PMTs, the quantity PMT_i is the i -th PMT input to the DNN in the form of charge information, and the weight w_i of the i -th PMT is a parameter in the DNN. When the training stage of the DNN ends, the w_i s would have reached their best values corresponding to a minima of the MSE . The PMT with the largest weight in magnitude indicates that the reconstruction

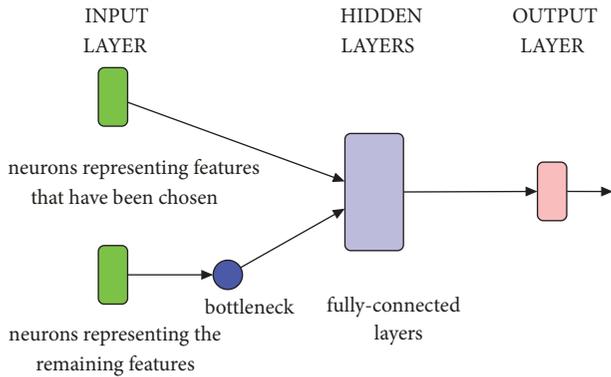


FIGURE 2: DNN architecture consisting of a bottleneck neuron.

of the position of the IBD prompt events relies the heaviest on this PMT compared to the rest of the candidate PMTs. Hence, this PMT would be our k -th important PMT, PMT_k^* . Crucially, this type of DNN only needs to be trained once to identify PMT_k^* no matter the value of k , whereas the brute force DNN needs to be trained $(N - k)$ times, once for each remaining candidate PMT.

4. Results

The heatmap in Figure 3 shows the resolutions as obtained from the residual distributions corresponding to using only one PMT for training and determining the vertex location of the antineutrino IBD interaction using the brute force search. Specifically, the resolution pertaining to using each PMT is indicated by a value from the color scale. Clearly, some PMTs contain more information about the vertex position than others. The first most important PMT from the brute force search, i.e., PMT_1^{brute*} , is chosen as the one having the smallest color value in the heatmap. The variation in resolution for the x -direction by column is due to the use of σ_x rather than σ_r . The reconstruction shows that the most important PMT is different for the x -direction and the z -direction. The heatmap pattern for the y -direction is similar to the x -direction, but with the dark region in the x -direction being the light region in the y -direction and vice versa, reflecting that x and y depend on $\cos \phi$ and $\sin \phi$, respectively in a cylindrical coordinate system, i.e., a $\pi/2$ shift in difference between x and y .

In Figure 4, the heatmap shows the weight corresponding to each PMT as obtained from the bottleneck neuron when searching for the first most important PMT. A higher weight indicates that the PMT has a larger impact on the vertex reconstruction. The PMT with the largest weight is identified as the most important PMT, PMT_1^{bneck*} . Ideally, we would like to constrain the weights to be discrete at the bottleneck, i.e., to be either 0 or 1, where weight is 1 for the most important PMT and 0 for the rest during the training of the DNN. However, such a constraint is nondifferentiable and noncontinuous with respect to the loss function which would render DNN parameter optimization using gradient

descent algorithms unfeasible. Comparing Figures 3 and 4, the brute force and the bottleneck DNN have chosen different PMTs as their most important PMTs possibly due to degeneracies in the detector. For example, in the z -direction, the PMTs at the top and bottom rings should produce the same resolution. Our suspect is that, during the training of the bottleneck DNN, some information sharing between a subset of PMTs, in which the DNN thinks their information values are similar, is unavoidable. Hence, the bottleneck neuron contains information from not one but a subset of PMTs; i.e., the importance by weights of each PMT could be partially “shared” amongst several PMTs. Further understanding of these are being conducted. Figures 5 and 6 are the results as obtained from the brute force and bottleneck DNN approach, respectively, while searching for the second most important PMT after having found the first most important PMT.

Figure 7 shows the residual curve for x and z as a function of the number of PMTs used in the reconstruction. Using an Nvidia Tesla P40 GPU, we estimated that it would take about 60 days to complete the entire residual curve with the brute force search, whereas it took about one day to complete with the bottleneck DNN. Resolutions from random choice of PMTs are also included in the figures as a comparison to the result from the bottleneck DNN. An empirical fit to the bottleneck DNN results is done with a triple exponential fit. It can be clearly seen that there is a diminishing return on the improvement in the vertex resolution when adding additional PMTs to an existing set of PMTs, which is an implication of the submodular [30] nature of the Gaussian standard deviation and its relationship to the information entropy, $H = \log \sigma + 0.5 \log(2\pi e)$. Succinctly, the submodularity of the Gaussian standard deviation, i.e., the resolution in this case shows that there is less new information that could be gained from adding a new PMT to a larger set of already-chosen PMTs than to a smaller set. As all the PMTs in the Daya Bay detector are of the same size and model, Figure 7 could also be interpreted as the residual curve being a function of the detector coverage.

5. Summary

In this work, we provide a fast approach using a deep neural network with a bottleneck neuron to uncover the effects of the number of photon sensors such as PMTs on the vertex resolutions in an antineutrino detector. The results have been compared with a random PMT search and a brute force search which yields the ideal result. Our inputs are the simulated charge information of the Daya Bay PMTs. The fast approach produces results close to those from the brute force search and fares much better than a random search. We find that the vertex resolution of the event reconstruction at the Daya Bay is approximately a multiexponentially decreasing function with respect to the number of PMTs and hence, also, the coverage. In future work, we envisage the possibility of incorporating the temporal information, i.e., the time of arrival of each photon in addition to the charge information to reconstruct the vertices. In addition, one could also study the size of the PMT needed alongside its installation location corresponding to the best event vertex reconstruction

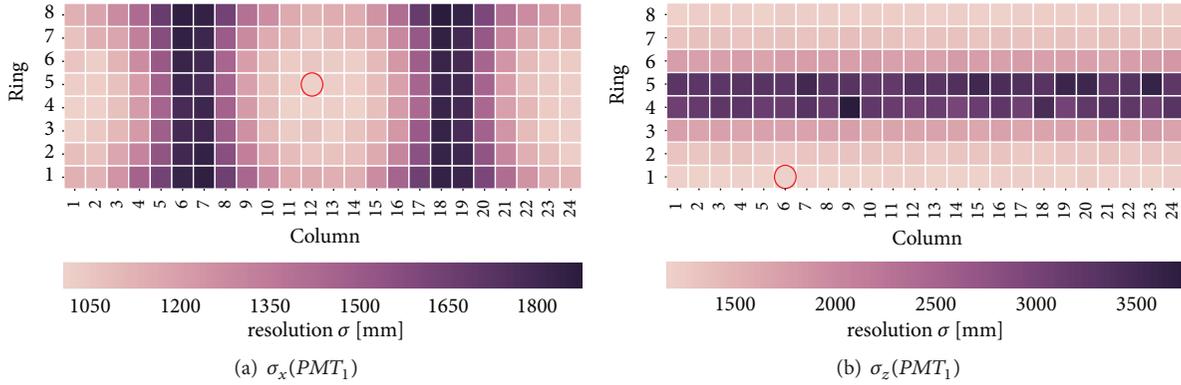


FIGURE 3: Resolution σ corresponding to using each individual PMT for (a) x and (b) z , using the brute force search. In each heatmap, the circled PMT corresponds to the best resolution amongst all the 192 PMTs.

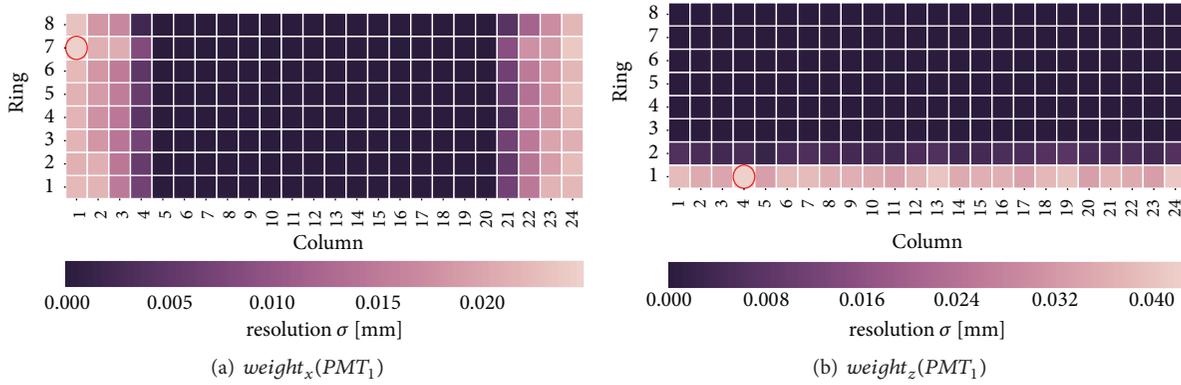


FIGURE 4: Weights as given by the bottleneck neuron corresponding to each PMT for (a) x and (b) z using the bottleneck DNN. In each heatmap, the circled PMT corresponds to the largest weight in the heatmap amongst all the 192 PMTs, indicating that it is the most heavily used PMT in the DNN during the vertex reconstruction.

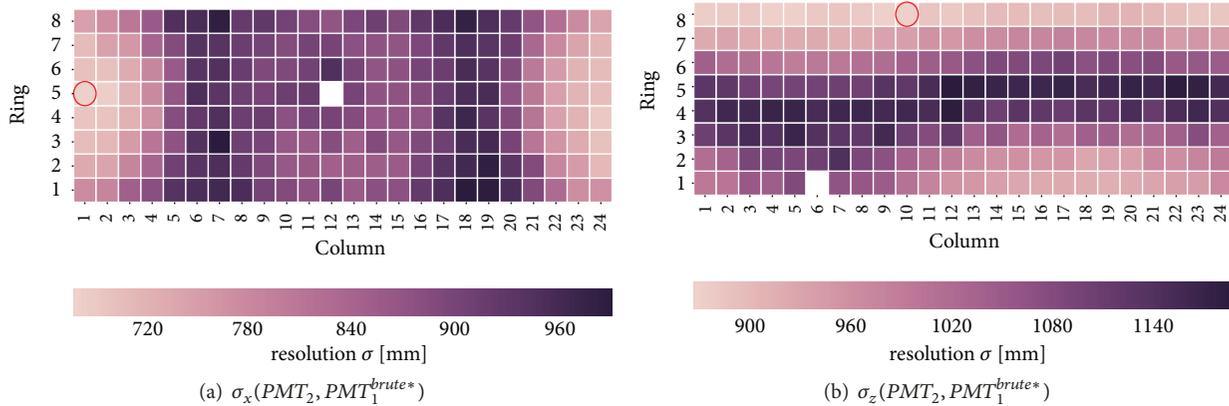


FIGURE 5: Resolution σ corresponding to using the charge information from the most important PMT as identified previously and a second candidate PMT for (a) x and (b) z . The first PMT as found by the previous round, PMT_1^{brute*} , using the brute force search is whitened in each figure. In each of the heatmap, the circled PMT, PMT_2^* , corresponds to the best resolution when considering the said PMT and the first important PMT found previously for the vertex reconstruction.

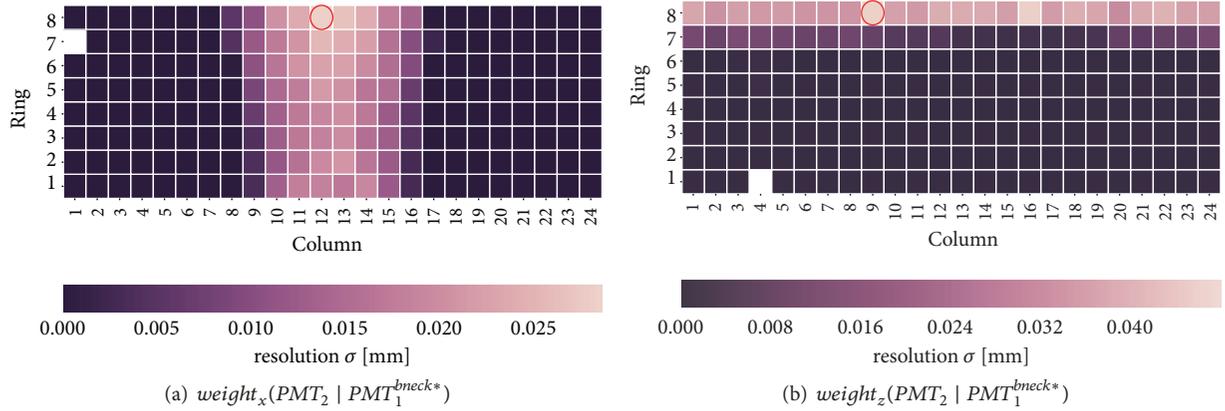


FIGURE 6: Weights as given by the bottleneck neuron corresponding to using a second candidate PMT, and the most important PMT, PMT_1^{bneck*} , as identified previously for (a) x and (b) z . PMT_1^{bneck*} is whitened in each figure. In each of the heatmap, the circled PMT, PMT_2^* , corresponds to the largest weight amongst the remaining 191 PMTs given that the first important PMT has been identified with the bottleneck DNN.

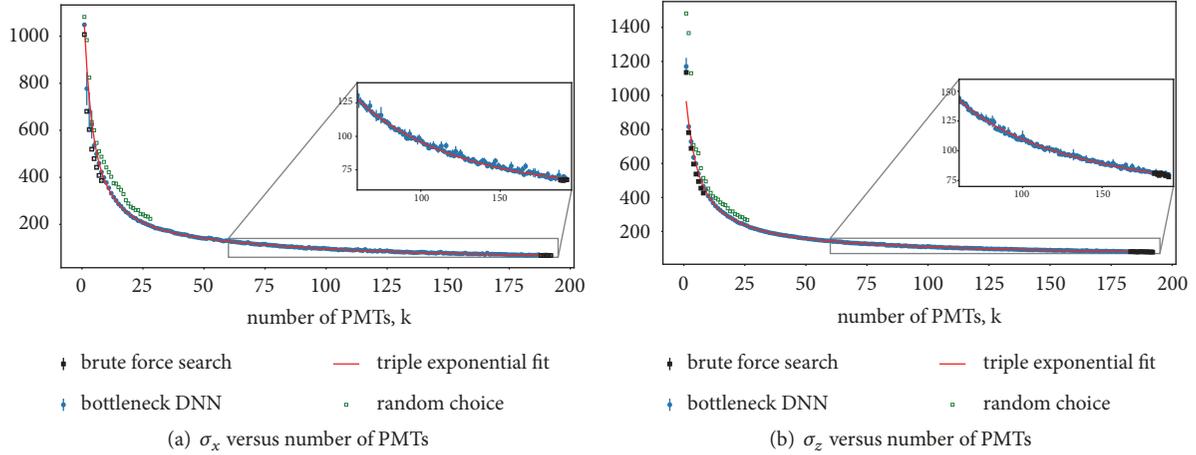


FIGURE 7: Residual curves for x (a) and z (b) as a function of the number of PMTs used in the reconstruction. The results plotted as solid squares are those from the brute force search. An empirical fit to the bottleneck DNN results was done with a triple exponential. Resolutions from random choice of PMTs are shown as empty squares. A zoom-in on the residual curves is also shown in the figures.

resolution. Also, a subsequent work from here would be the study of the effect of PMTs based on the event energy upon obtaining the vertex positions. Although studying the energy might need modifications to the deep network as the energy is a positive-definite quantity, the energy resolution is important when considering physics sensitivity and thereby also impacting the design of future antineutrino detectors including JUNO.

In order to use the bottleneck DNN approach for new detectors in designing phases, we suggest Monte Carlo simulations using various N . Then, one can obtain the (P, N, k) surface in the hyperspace, where P is some detector performance metric; N and k are as described in this work. Experimentalists can then decide on the (P, N, k) working point for their detector in accordance with their construction budget and the desired detector performance.

Data Availability

The data used to support the findings of this study are available from the corresponding author upon request.

Disclosure

Zhi-Qiang Qian is co-first author.

Conflicts of Interest

The authors declare that there are no conflicts of interest regarding the publication of this paper.

Acknowledgments

The authors thank Shen-Jian Chen and Zuo-Wei Liu for their computing facilities and helpful discussions. They would

also like to thank Chao Zhang, Zhe Wang, Samuel Kohn, the Daya Bay ACC, and the Collaboration for their time and comments. This work was supported by the National 973 Project Foundation of the Ministry of Science and Technology of China (Contract no. 2013CB834300) and the International Science & Technology Cooperation Program of China (Contract No. 2015DFG02100).

References

- [1] J. Cao and K.-B. Luk, “An overview of the Daya Bay reactor neutrino experiment,” *Nuclear Physics B*, vol. 908, pp. 62–73, 2016.
- [2] Double Chooz collaboration, Y. Abe et al., “Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector,” *Journal of High Energy Physics*, vol. 2014, no. 86, 2014.
- [3] RENO Collaboration, J. Choi et al., “Observation of Energy and Baseline Dependent Reactor Antineutrino Disappearance in the RENO Experiment,” *Physical Review Letters*, vol. 116, 2016.
- [4] JUNO collaboration, F. An et al., “Neutrino physics with JUNO,” *Journal of Physics G: Nuclear and Particle Physics*, vol. 43, Article ID 030401, 2016.
- [5] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [6] Y. Bengio, “Deep Learning of Representations: Looking Forward,” in *Statistical Language and Speech Processing*, A. Dediu, C. Martín-Vide, R. Mitkov, and B. Truthe, Eds., p. 37, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [7] J. Schmidhuber, “Deep learning in neural networks: an overview,” *Neural Networks*, vol. 61, pp. 85–117, 2015.
- [8] A. Krause, J. Leskovec, C. Guestrin, J. VanBriesen, and C. Faloutsos, “Efficient sensor placement optimization for securing large water distribution networks,” *Journal of Water Resources Planning and Management*, vol. 134, no. 6, pp. 516–526, 2008.
- [9] K. Worden and A. P. Burrows, “Optimal sensor placement for fault detection,” *Engineering Structures*, vol. 23, no. 8, pp. 885–901, 2001.
- [10] S. Jetter, D. Dwyer, W.-Q. Jiang et al., “PMT waveform modeling at the Daya Bay experiment,” *Chinese Physics C*, vol. 36, no. 8, pp. 733–741, 2012.
- [11] W. S. McCulloch and W. Pitts, “A logical calculus of the ideas immanent in nervous activity,” *Bulletin of Mathematical Biology*, vol. 5, no. 4, pp. 115–133, 1943.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds., vol. 25, pp. 1097–1105, 2012.
- [13] D. Silver, A. Huang, C. J. Maddison et al., “Mastering the game of Go with deep neural networks and tree search,” *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [14] F. Buggenthin, F. Buettner, P. S. Hoppe et al., “Prospective identification of hematopoietic lineage choice by deep learning,” *Nature Methods*, vol. 14, no. 4, pp. 403–406, 2017.
- [15] H. W. Lin, M. Tegmark, and D. Rolnick, “Why does deep and cheap learning work so well?” *Journal of Statistical Physics*, vol. 168, no. 6, pp. 1223–1247, 2017.
- [16] D. Guest, J. Collado, P. Baldi, S. Hsu, G. Urban, and D. Whiteson, “Jet flavor classification in high-energy physics with deep neural networks,” *Physical Review D: Particles, Fields, Gravitation and Cosmology*, vol. 94, no. 11, 2016.
- [17] P. T. Komiske, E. M. Metodiev, and M. D. Schwartz, “Deep learning in color: towards automated quark/gluon jet discrimination,” *Journal of High Energy Physics*, vol. 2017, no. 1, 2017.
- [18] P. Baldi, P. Sadowski, and D. Whiteson, “Searching for exotic particles in high-energy physics with deep learning,” *Nature Communications*, vol. 5, 2014.
- [19] P. Baldi, P. Sadowski, and D. Whiteson, “Enhanced Higgs Boson to Tau+ Tau- Search with Deep Learning,” *Physical Review Letters*, vol. 114, no. 11, 2015.
- [20] S. Kohn, “Understanding Backgrounds using Deep Learning at the Daya Bay Experiment,” in *Proceedings of the Work presented at the 28th International Symposium on Lepton Photon Interactions at High Energies*, 2017.
- [21] MiniBoone collaboration, R. Acciarri et al., “Convolutional neural networks applied to neutrino events in a liquid argon time projection chamber,” *Journal of Instrumentation*, vol. 12, 2017.
- [22] A. Aurisano, A. Radovic, D. Rocco et al., “A convolutional neural network neutrino event classifier,” *Journal of Instrumentation*, vol. 11, no. 09, pp. P09001–P09001, 2016.
- [23] NEXT collaboration, J. Renner et al., “Background rejection in NEXT using deep neural networks,” *Journal of Instrumentation*, vol. 12, Article ID T01004, 2017.
- [24] E. Racah, S. Ko, P. Sadowski et al., “Revealing fundamental physics from the Daya Bay Neutrino Experiment using deep neural networks,” in *Proceedings of the 15th IEEE International Conference on Machine Learning and Applications, ICMLA 2016*, pp. 892–897, USA, December 2016.
- [25] L. Prechelt, “Early Stopping - But When?” in *Neural Networks: Tricks of the Trade*, vol. 1524 of *Lecture Notes in Computer Science*, pp. 55–69, Springer Berlin Heidelberg, Berlin, Heidelberg, 1998.
- [26] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *Nature*, vol. 323, no. 6088, pp. 533–536, 1986.
- [27] V. Nair and G. Hinton, “Rectified Linear Units Improve Restricted Boltzmann Machines,” in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, J. Furnkranz and T. Joachims, Eds., pp. 807–814, Omnipress, 2010.
- [28] J. Bergstra, R. Bardenet, B. Yoshua, and K. Balázs, “Algorithms for Hyper-parameter Optimization,” in *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS11*, pp. 2546–2554, Curran Associates Inc, USA, 2011.
- [29] G. E. Hinton and R. S. Zemel, “Autoencoders, Minimum Description Length and Helmholtz Free Energy,” in *Advances in Neural Information Processing Systems 6*, J. D. Cowan, G. Tesauro, and J. Alspector, Eds., pp. 3–10, Morgan-Kaufmann, 1994.
- [30] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, “An analysis of approximations for maximizing submodular set functions—I,” *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, 1978.



Hindawi

Submit your manuscripts at
www.hindawi.com

