



OPEN ACCESS

EDITED BY

Andrey Ustyuzhanin,
National Research University Higher
School of Economics, Russia

REVIEWED BY

Jonathan Andrew Miller,
Onto Innovation, United States
Stefano Belforte,
National Institute of Nuclear Physics
of Trieste, Italy

*CORRESPONDENCE

Blair Jamieson
bljamieson@uwinnipeg.ca

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 26 June 2022

ACCEPTED 29 August 2022

PUBLISHED 30 September 2022

CITATION

Jamieson B, Stubbs M, Ramanna S,
Walker J, Prouse N, Akutsu R, de
Perio P and Fedorko W (2022) Using
machine learning to improve neutron
identification in water Cherenkov
detectors. *Front. Big Data* 5:978857.
doi: 10.3389/fdata.2022.978857

COPYRIGHT

© 2022 Jamieson, Stubbs, Ramanna,
Walker, Prouse, Akutsu, de Perio and
Fedorko. 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.

Using machine learning to improve neutron identification in water Cherenkov detectors

Blair Jamieson^{1*}, Matt Stubbs², Sheela Ramanna²,
John Walker^{1,3}, Nick Prouse³, Ryosuke Akutsu³,
Patrick de Perio^{3,4} and Wojciech Fedorko³

¹Physics Department, University of Winnipeg, Winnipeg, MB, Canada, ²Applied Computer Science Department, University of Winnipeg, Winnipeg, MB, Canada, ³Science Division, TRIUMF, Vancouver, BC, Canada, ⁴Kavli IPMU (WPI), UTIAS, The University of Tokyo, Tokyo, Japan

Water Cherenkov detectors like Super-Kamiokande, and the next generation Hyper-Kamiokande are adding gadolinium to their water to improve the detection of neutrons. By detecting neutrons in addition to the leptons in neutrino interactions, an improved separation between neutrino and anti-neutrinos, and reduced backgrounds for proton decay searches can be expected. The neutron signal itself is still small and can be confused with muon spallation and other background sources. In this paper, machine learning techniques are employed to optimize the neutron capture detection capability in the new intermediate water Cherenkov detector (IWCD) for Hyper-K. In particular, boosted decision tree (XGBoost), graph convolutional network (GCN), and dynamic graph convolutional neural network (DGCNN) models are developed and benchmarked against a statistical likelihood-based approach, achieving up to a 10% increase in classification accuracy. Characteristic features are also engineered from the datasets and analyzed using SHAP (SHapley Additive exPlanations) to provide insight into the pivotal factors influencing event type outcomes. The dataset used in this research consisted of roughly 1.6 million simulated particle gun events, divided nearly evenly between neutron capture and a background electron source. The current samples used for training are representative only, and more realistic samples will need to be made for the analyses of real data. The current class split is 50/50, but there is expected to be a difference between the classes in the real experiment, and one might consider using resampling techniques to address the issue of serious imbalances in the class distribution in real data if necessary.

KEYWORDS

machine learning, graph neural networks, water Cherenkov detector, particle physics, neutrino physics

1. Introduction

One exciting frontier within experimental neutrino physics is the improved identification of neutrons from inverse beta decay reactions ($\nu_e + p^+ \rightarrow e^+ + n$). This task, referred to as “neutron tagging,” is particularly challenging due to the low energy scale and faint signals involved. Progress in this field could lead to a host of advancements

in particle physics, including a first detection of diffuse supernova background neutrinos (Fernández, 2016), and improved understanding of the neutrino mass hierarchy and the CP violating phase (Irvine, 2014). However, Water Cherenkov (WC) detectors have historically been limited in their detection capability of these low energy neutron capture events.

Neutrons are commonly liberated in water due to the inverse beta decay (IBD) process, in which an electron antineutrino collides with a proton to yield a positron and a free neutron. From there, the free neutron undergoes thermalization, colliding with neighboring molecules and gradually losing energy until it reaches water temperature. Approximately 200 μ s after thermalization, the free neutron is captured by a proton or oxygen nucleus, releasing a gamma particle γ at 2.2 MeV ($n + p \rightarrow d + \gamma$) (Watanabe et al., 2009) where d is deuterium (or “heavy hydrogen”), the isotope of hydrogen with a proton and neutron in the nucleus. The capture cross-section of this neutron capture on a hydrogen nucleus (proton) is only 0.33 barns, and the resulting 2.2 MeV gamma produces such a faint light signal that it is very difficult to identify by the Photomultiplier Tubes (PMTs) in a WC detector. The detection of the signal gamma-ray produced by the neutron capture relies on the detection of Cherenkov photons produced by Compton scattered electrons produced by the gamma-ray. Many traditional WC detectors actually have thresholds of 5 MeV, high enough that none of these signals would be recorded at all.

To address this problem, the addition of gadolinium chloride (GdCl_3 ; a light, water soluble-compound) to the SK detector water was proposed in 2003 (Beacom and Vagins, 2004). Gadolinium is known for having the “largest capture cross-section for thermal neutrons among all stable elements” (Ankowski et al., 2012). At $\sim 49,700$ barns, the gadolinium capture cross-section is over six orders of magnitude larger than for free protons, leading to faster captures. Neutron capture on gadolinium also leads to an 8 MeV cascade of gammas (7.9 MeV cascade 80.5% of the time and an 8.5 MeV cascade 19.3% of the time; Watanabe et al., 2009), a signal which is far easier to detect due to its relatively higher energy. Beacom and Vagins showed that only a 0.1% addition of gadolinium by mass leads to at least a 90% probability of neutron capture on gadolinium (the other 10% or less of neutron captures are still by hydrogen nuclei). In addition, the neutron capture by gadolinium after thermalization occurs in roughly 20 μ s, nearly 10 times more quickly than capture on protons.

This paper presents the implementation of several machine learning methods that attempt to improve the efficiency of neutron tagging for simulations of neutron capture and background radiative neutrino events within the gadolinium-doped intermediate WC detector (IWCD) for Hyper-K (Proto-Collaboration et al., 2018). Since the machine learning methods are fast once the training is completed, they can be used for the semi-offline analysis soon after data is taken to monitor neutron detection rates. This could be important to monitor event rates

and understand if there are any backgrounds that are changing in the detector when it is running. The methods may also get used in later stages of the off-line analysis, particularly if they outperform more traditional cut-based methods.

The structure of the paper is as follows. Section 2 discusses related works in the intersecting fields of particle physics, neutron tagging and machine learning. Section 3 then introduces the relevant machine learning theories and algorithms used in this research, including boosted decision trees (XGBoost), SHAP (SHapley Additive exPlanations), and graph neural networks (GNNs). In Section 4, the data and data simulation process are explored. Also in this section, a likelihood analysis benchmark is shown based on event hit totals and charge sums. Section 5 illustrates the process of engineering characteristic features from the data and covers the implementation and tuning of the XGBoost model. Afterward, an analysis of relative feature importances is applied using SHAP. Section 6 presents the results of the GCN and DGCNN graph neural network models and discusses various methods of graph network construction. One of the main goals of this research is to investigate the applicability, performance and feasibility of GNNs on the IWCD particle data, in particular for the low energy regime where the number of event hits is small and CNNs tend to struggle. Finally, Section 7 concludes on the findings of the previous chapters.

2. Related work

2.1. Machine learning in particle physics

The uses of machine learning and its historical development in the field of particle physics is discussed in Bourilkov (2019). Traditional means of event selection in particle physics are discussed in both Bourilkov (2019) and Guest et al. (2018). These methods often involved a series of boolean “cuts” (decisions) on single variables at a time, followed by statistical analyses on the remaining data. However, over the past several decades, physicists have developed algorithms that employ machine learning to study multiple variables simultaneously in multivariate analysis (MVA). Guest et al. (2018) describes the use of an assortment of machine learning techniques for MVA in the physics context, include support vector machines, kernel density estimation, random forests, boosted decision trees, etc. Carleo et al. (2019) provides an overview of applications of machine learning within the physical sciences, including applications to quantum computing, chemistry, and cosmology. Carleo et al. (2019) also discusses applications to particle physics, including jet physics and neutrino signal classification. Machine learning applications are discussed for a variety of neutrino experiments, including the MicroBooNE collaboration, Deep Underground Neutrino Experiment (DUNE) and the IceCube Observatory at the South Pole.

2.2. Boosted decision trees

Boosted decision trees (BDTs) are a commonly applied machine learning method in modern particle physics analysis. For example, [Roe et al. \(2005\)](#) details the improved performance of particle classification in the MiniBooNE experiment, which searches for neutrino oscillations, using BDTs compared to artificial neural networks. [Radovic et al. \(2018\)](#) discusses multiple use cases of BDTs at the Large Hadron Collider (LHC) at CERN, including the application of BDTs to improve the energy reconstruction (mass resolution) of the CMS (Compact Muon Solenoid) calorimeter, as well as the implementation of BDTs to improve the sensitivity of the ATLAS detector to various Higgs boson decay modes. For the latter, the sensitivity of diphoton decay ($H \rightarrow \gamma\gamma$) and antitau-tau pair decay ($H \rightarrow \tau^+\tau^-$) were improved by an amount equivalent to adding 50 and 85% more data to the detector, respectively. Beyond learning tasks, BDTs can also be used at the early stages of the machine learning lifecycle. For example, [Gligorov and Williams \(2013\)](#) modifies the standard boosted decision tree algorithm to improve high-level triggering in detector data acquisition systems. A general BDT usage guidebook is presented in [Cornell et al. \(2022\)](#) for the hypothetical identification of the smuon particle and performance is compared to the classic “cut-and-count” approach.

2.3. Deep learning and graph neural networks

The computer vision approach to particle classification, which consists of reconstructing particle events as images and applying convolutional neural networks (CNNs), has been applied in various detector experiments ([Macaluso and Shih, 2018](#); [ATLAS Collaboration, 2019](#); [Andrews et al., 2020](#)). However, the conversion of data from irregular detector geometries into a two-dimensional grid for images inherently causes loss of information. For events with few hits, the sparsity of the resulting image is also difficult for CNNs to learn from (e.g., [Shlomi et al., 2021](#)). Alternately, deep learning sequence models, inspired by tasks in natural language processing, have also been adapted to the particle physics domain by modeling particles and measurement objects in a sequential order. Instances of this approach include tagging of jets containing b -hadrons in the ATLAS experiment ([ATLAS Collaboration, 2017](#)) and classifying energetic hadronic decays in the CMS experiment ([Sirunyan et al., 2020](#)). However, the imposed ordering of objects in the sequence constrains the learning of the model. The limitations of both computer vision and sequence deep learning approaches are discussed in [Shlomi et al. \(2021\)](#).

Graph neural networks (GNNs) represent an emerging architectural class of deep learning which undertakes to learn from data structured in a graph format, for which particle events find a natural representation. [Shlomi et al. \(2021\)](#) surveys the theory and applications of GNNs in particle physics. The graph classification task is partitioned into jet classification and event classification. While jets represent a part of a particle collision occurrence, an event references the full history of the particular physics process. In [Qu and Gouskos \(2020\)](#), the jet is viewed as an unordered structure of particles, analogous to the point cloud representation of shapes in 3D space. The authors propose the “ParticleNet” method, which uses the “EdgeConv” block as an analog for CNN convolution on 3D point clouds and updates the graph representation dynamically, and report state-of-the-art performance on jet tagging tasks. For event classification, one example is the deployment of GNNs in the IceCube neutrino observatory ([Choma et al., 2018](#)). In this case, the irregular hexagonal geometry of the detector is itself modeled as a graph, where the sensors are the graph nodes and the edges represent their connections. Given the sparsity of activated sensors in an event, every event is considered as a different graph composed only of the active sensors in the event. Although learning occurs over relatively small sample sizes, the authors report an approximate 3x improvement in signal-to-noise ratio compared to the physics baseline and the CNN approach.

3. Machine learning methods studied

3.1. XGBoost

Over the last several years, the machine learning model “XGBoost” has gained popularity for its performance in classification or regression tasks involving tabular data over a variety of domains, including vehicle accident detection ([Parsa et al., 2020](#)), cancer diagnostics ([Tahmassebi et al., 2019](#)), network intrusion detection ([Bhattacharya et al., 2020](#)) and Higgs boson identification ([Chen and He, 2015](#)). XGBoost stands for “eXtreme Gradient Boosting.” In general, gradient boosting refers to the process of beginning with a single weak learner and iteratively constructing superior learners that improve on the errors of their predecessors. The new learners attempt to optimize an overall loss function over the problem space by each following the negative gradient of the loss function.

XGBoost was introduced by [Chen and Guestrin \(2016\)](#) in their paper, which considered the case of decision trees as the individual learners in the function ensemble. In general, a decision tree applies classification or regression to an example by partitioning the example through a series of splits (decisions) from the root node to a leaf of the tree. The given tree splits are themselves computed by calculating which partition leads to maximum information gain. For any specific training example,

the overall output is the additive sum of the outputs from every individual tree. To apply gradient boosting in the context of decision trees, an appropriate objective function (loss) must be defined. Chen and Guestrin define the overall objective function as the sum of a regular loss and a regularization term. Practically, when constructing a given decision tree in the XGBoost ensemble, it is too computationally expensive to iterate through all possible tree structures and compute the objective function for each possibility. Instead, a greedy approach is applied where, starting at the tree node, branches are successively added by finding the particular split which leads to maximum gain.

3.2. SHAP

The Shapley value (which SHAP derives from) traces back to Lloyd Shapley's paper "stochastic games," published in Princeton in 1953 (Shapley, 1953). At the time, Shapley was studying the field of cooperative game theory and searching for a mapping from a coalition single game to a numeric payoff vector. Shapley found an intuitive solution to the seemingly intractable problem by searching for a set of "reasonable axioms" (efficiency, symmetry, dummy, and additivity; Shapley, 1953). His resulting "Shapley value" can be viewed as an "index for measuring the power of players in a game" (Winter, 2002). In the context of physical event classification, the player is analogous to the event feature, the game is analogous to the event and the label is the analogous to the numeric payoff output.

Winter's paper (Winter, 2002) reviews the theoretical framework for the derivation of the Shapley values. Lundberg and Lee (2017) extend this definition, introducing the "SHAP" values as the Shapley values of a "conditional expectation function of the original model." They also present the concept of the "explanation model" in which the output prediction of the ML model may be viewed as a model itself. Their definition of an "Additive Feature Attribution Method" is one in which the explanation model may be represented as a linear function of binary variables. This makes it possible to view the marginal contributions of individual features for any given event.

3.3. Graph neural network (GNN)

While traditional machine learning algorithms have proven effective at learning from tabular data, they have historically struggled to learn well from natural data, including images, natural language or audio. While deep learning architectures like convolutional neural networks (CNN; Krizhevsky et al., 2017) and recurrent neural networks (RNN; Mikolov et al., 2010) have proven effective at learning from image or sequence data, *geometric deep learning*, the umbrella term for the task of deep learning on graph data, is an emerging area of research. Where a

given graph G may be denoted by its set of vertices and edges $G = \{V, E\}$, the nodes represent objects or concepts and the edges represent their relationships. A variety of situations may be modeled by graphs, including social networks, molecules, Internet traffic, etc. (Zhou et al., 2020). The GNN is designed to operate directly on data input as a graph. Low energy neutrino-induced events in the IWCD may be naturally represented by a graph, where the PMTs constitute the nodes and the edges represent the connections between the PMTs.

The origin of deep learning on graphs traces back to the late 1990s, when RNNs were applied to directed, acyclic graphs (directional edges, no loops formed by a collection of edges; Zhou et al., 2020). Using this approach, node feature states are updated in successive layers until equilibrium is reached. This technique was later generalized to cyclic graphs as well in 2008 (Scarselli et al., 2008). Soon after, following the widespread success of CNNs, significant interest grew in generalizing some concepts from CNNs to learning on graphs. The first successful adaption of the convolution operation to graphs was developed by Bruna et al. (2013) in 2013 using Laplacian eigenvectors. The computational complexity of this procedure was later greatly reduced by applying polynomial spectral filters instead of Laplacian eigenvectors (Michael et al., 2016; Kipf and Welling, 2017). Approaches have also been developed which apply spatial, and not spectral, filters for the convolutional operation (Monti et al., 2017). In general, GNNs apply a series of filtering and activation layers to update the feature representation of every node. Once the network has passed all the hidden layers, the output node labels may be used directly in node-focused tasks, or the node outputs may be pooled together to obtain an overall coarsened representation for graph classification.

3.3.1. Graph convolutional network

Kipf and Welling demonstrated the successful approach of using a convolutional architecture to learn on graphs in their paper "Semi-supervised classification with graph convolutional networks" (Kipf and Welling, 2017). This approach applies an approximation of spectral graph convolution. The spectral decomposition of a graph denotes the breakdown of the graph's Laplacian matrix \mathcal{L} into its elementary orthogonal components, i.e., the *eigendecomposition* of \mathcal{L} . The graph Laplacian \mathcal{L} represents a graph in matrix format and is a graphical analog to the familiar Laplacian operator for multivariate and continuous functions. For a graph $G = \{V, E\}$, $\mathcal{L}(G)$ is equal to the difference between the degree matrix D (diagonal matrix where every element represents the degree, i.e., number of connections of the corresponding vertex) and adjacency matrix A (matrix with vertices labeled by rows and columns where 0 and 1 s represent nonadjacent and adjacent pairs of vertices) of G . However, the computation of \mathcal{L} is computationally expensive and can be a procedural bottleneck. Hammond et al. (2011) proposed a computation of \mathcal{L} using the first K Chebyshev polynomials

that avoids diagonalization. By taking the first-order Chebyshev approximation $K = 1$ and further constraining other parameters, the multi-layer GCN propagation rule is reached,

$$H^{l+1} = \sigma(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^l W^l), \quad (1)$$

where H^l and H^{l+1} denote the node feature matrices at layers l and $l + 1$, $\tilde{A} = A$ (adjacency matrix of graph) + I_N (identity matrix), $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$, W^l denotes the matrix of weights at layer l and σ is an activation function such as the rectified linear activation unit (ReLU).

3.3.2. Dynamic graph convolutional neural network

The dynamic graph convolution neural network (DGCNN), introduced by Wang et al. (2019), was designed specifically to learn from point cloud graphs for segmentation or classification tasks. Point clouds are collections of three-dimensional coordinates (points) in Euclidean space. However, the DGCNN model also allows the graph nodes to include other features in addition to the spatial coordinates. The main feature of the DGCNN model is the introduction of the “EdgeConv” convolutional operator. EdgeConv is designed to learn edge features between node pairs, i.e., a node and its neighboring connections. The DGCNN model is dynamic because, for every EdgeConv block, the graph representation is updated. This departs from the action of operating on a fixed graph like most other GNN architectures.

In the DGCNN model, a series of EdgeConv layers are applied to the graph. For a given layer in the network, the EdgeConv operation is applied for every node and its k nearest neighbors in semantic space, where k is a tunable hyperparameter. For two neighboring nodes \mathbf{x}_i and \mathbf{x}_j , a fully connected layer $h_{\Theta}()$ with learnable weights Θ and an adjustable number of compute units is applied to learn the pairwise edge features \mathbf{e}_{ij} . The node representations are then updated by aggregating these edge features over the node neighborhood. The EdgeConv filter $h_{\Theta}(\mathbf{x}_i, \mathbf{x}_j) = h(\mathbf{x}_i, \mathbf{x}_j - \mathbf{x}_i)$ operates over individual nodes and local node neighborhoods, thus allowing the model to learn both local neighborhood structure and global graph structure. In addition, the dynamic recomputation of the graph for every EdgeConv layer allows for groupings of nodes in semantic space compared to the fixed spatial input space, allowing for a diffusion of information throughout the entire graph.

4. Data analysis

4.1. Data simulation

The data used in this research was simulated using WCSim software to generate neutron and background electron events

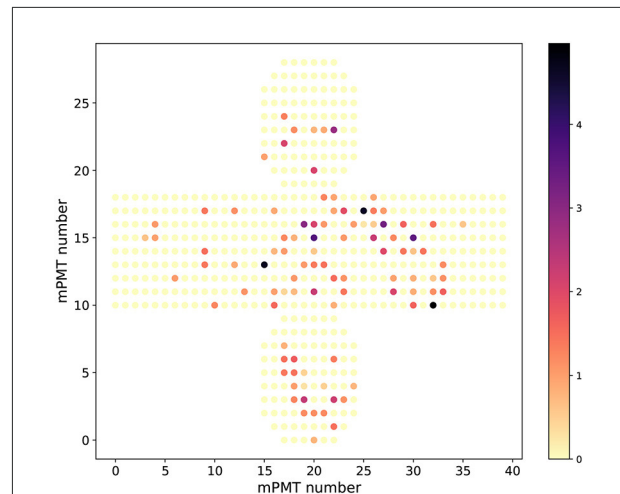


FIGURE 1
Unrolled cylinder event displays showing the charged deposit in units of photoelectrons as the colored points for a sample electron background event. Multi PMT modules without any charge are shown in yellow.

for the IWCD detector. WCSim, designed to recreate physics events within large WC detectors (O’Sullivan, 2021), is based on Geant4 (Agostinelli, 2003) and also depends on ROOT (Brun and Rademakers, 1997). The simulations used a cylindrical tank with a height of 6 m and a diameter of 8 m, and with 525 multi-PMT (mPMT) modules of 19 Hamamatsu PMTs each lining the walls of the simulated detector. With a PMT dark noise rate of 1 kHz and gadolinium doping of 0.1% by mass in the water to generate an approximate 90% thermal neutron capture on gadolinium nuclei, the simulations procured datasets of about 1.6 million events in total divided nearly evenly between neutron capture and background electron events. The current samples used for training are representative only, and more realistic samples will need to be made for the analyses of real data. The current class split is 50/50, but there is expected to be a difference between the classes in the real experiment, and one might consider using resampling techniques to address the issue of serious imbalances in the class distribution in real data if necessary.

The data was saved in a three-dimensional format of (event, hit, features) where the eight feature values stored were the charge, time, 3D position (x, y, z) measured relative to the center of the cylindrical shape of the detector and 3D orientation (dx, dy, dz) of each hit PMT. The z is along the direction of the beam, y is vertical, and x is chosen to maintain a right-handed coordinate system. The detector studied is cylindrical, and an event display mapping the PMT locations on the cylinder to a flat image are shown in Figure 1. In this simulation, the detectors consist of modules of 19 PMTs, and therefore several modules may have multiple photoelectrons, but this is a summary display showing the total over the 19 PMTs in each module.

Other features may be engineered from these base eight, a topic which is explored in Section 5.

Multiple datasets were tested with electron background distributions of different uniform energy levels (i.e., a uniform 8 or 20 MeV energy distribution). Comparisons with the different background distributions can be found in the thesis of [Stubbs \(2021\)](#). However, to generate a more realistic approximation of the background in the IWCD, an electron (beta decay) background was simulated following the energy spectrum of decays of isotopes produced by cosmic ray muon spallation. Only this more realistic background is used in this paper. In her presentation on muon spallation background in the Super-Kamiokande experiment, [Bernard](#) notes that at lower energy scales (tens of MeVs), muon spallation is a dominant source of background ([Bernard, 2019](#)). Due to the high muon flux at sea level of $6.0 \times 10^5 \text{ m}^{-2} \text{ hr}^{-1}$ ([Li and Beacom, 2014](#)), SK was built under 1,000 m of rock. The muons lose energy as they travel through the rock, leading to a far reduced flux rate of $9.6 \text{ m}^{-2} \text{ hr}^{-1}$ at the detector. The IWCD, however, is to be deployed in only a 50 m deep pit. Therefore, the spallation flux will be greater for IWCD, and it is even more important to reduce this background for identifying neutron captures at low energies. The combined muon spallation energy spectra from [Bernard](#) was used as an input to WCSim, replicating the SK spallation energy distribution for the simulation of electron background radiation events in the IWCD detector. The resulting electron background energy distribution follows a right-skewed distribution from ~ 0 to 16 MeV. This background, along with the regular neutron capture events generated by WCSim, constituted the dataset used in this research.

4.2. Likelihood baseline analysis

As shown in [Figures 2A,B](#), the difference in the total number of hits and charge sums between neutron and background electron events is the most obvious source of separability between these event types (the rest of the distributions in [Figure 2](#) will be discussed in the following sections). For the low energy events being considered, there is close to a 100% correlation between these variables, since each PMT hit is most likely a single photon hit, and only occasionally two photons. A statistical likelihood analysis based on these features was implemented to determine a baseline classification accuracy, defined as the number of correct predictions divided by the total number of predictions, for later comparison against other machine learning approaches. The likelihood baseline classification accuracy was calculated by estimating the probability density function (PDF) of the neutron and electron events based on their nhit distributions and then classifying the events based on highest likelihood. The kernel density estimate (KDE) was used as an estimate of the underlying PDF for the corresponding distribution. The density of the KDE instance,

once fit over a distribution of data, was then used to evaluate the event likelihood at a given point.

Univariate KDEs were calculated for the neutron and electron events based on their “nhits” distributions on training events. The final evaluation type, “nhits,” involved calculation of KDEs for neutron and electron events on the training set for the distribution of number of hits. All events in the test set were then classified based on the highest density of the neutron and electron multivariate KDEs.

The likelihood classification approach using univariate KDEs yielded a classification accuracy of 62.4%. The runtime cost of classifying events from highest KDE likelihood was ~ 1 h and 20 min on average for the testing set, while fitting the univariate KDEs to the training dataset only took a few minutes.

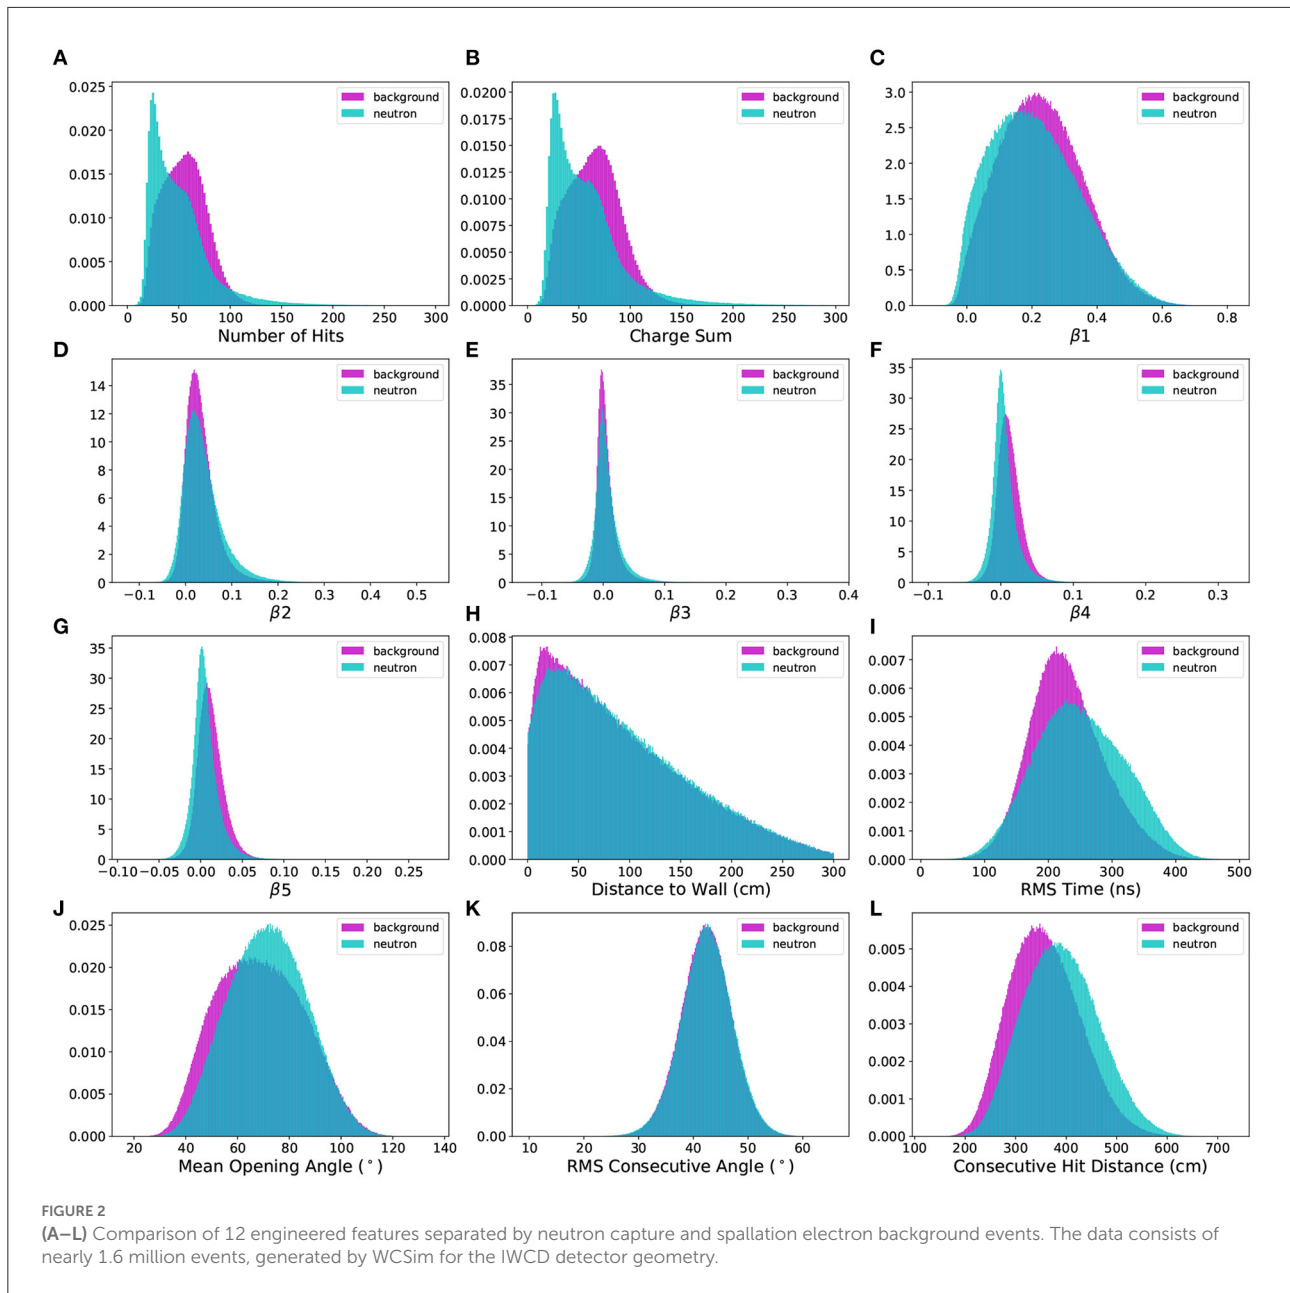
5. Feature engineering

In machine learning, feature engineering is the process of applying domain knowledge to extract useful features from the original dataset. These features are often more useful than the raw data itself for predictive or analytic tasks. However, the features must be carefully selected to extract as much information from the data as possible. Thus, a search was conducted for useful features in the domain of neutron capture in WC detectors. Relevant features were selected to aggregate information from each event, reducing the complexity of the dataset and extracting it into a more useful format. It was found that the classification performance of the XGBoost models significantly improved upon application to the aggregated features compared to the original dataset.

5.1. Beta parameters

One way to quantify event topology is by the amount of anisotropy within the event with respect to the event vertex (the vertex position denotes the Cartesian coordinates of the start of the event). For comparison of neutron capture to background events, isotropy may be a discriminating factor due to the backgrounds being single electron events, while the neutron signal is multiple gammas from a neutron capture. Several isotropy parameters were considered for use in this study, including Θ_{ij} , “the average of the angles between each pair of PMT hits in an event with respect to the fitted vertex position,” the correlation function ring inner product (CFRIP), which compares the angular correlation of the event to that of a perfect ring, and the beta parameters $\beta(l)$, defined similarly to Θ_{ij} but which make use of Legendre polynomials ([Wilson, 2015](#)).

Both [Wilson \(2015\)](#) and [Dunmore \(2004\)](#) found the beta parameters to yield the most powerful discrimination based on event isotropy between different types of subatomic particle events. Following this result, the beta parameters were chosen as



the measure of isotropy in this project. The definition for the l -th beta parameter $\beta(l)$ is

$$\beta(l) = \langle P(l)(\cos \theta_{ik}) \rangle_{i \neq k}, \quad (2)$$

where $\beta(l)$ is equal to the average of the l -th Legendre polynomial $P(l)$ of the cosine of the angle θ_{ik} between every pair of hit PMTs in the event ($i \neq k$) with respect to the event vertex. For any of the beta parameters, a value of 0 indicates perfect isotropy, while higher absolute values indicate directionality and lower isotropy. The beta parameter distributions for the datasets in this paper are shown in Figures 2C–G for β_1 through

β_5 , respectively. In practice, an event vertex would need to be calculated using an existing vertex reconstruction method. For the purpose of this study, the truth information is used for the exact event vertex position.

5.2. Time of flight

The root-mean-square (RMS) time of flight was selected as an engineered feature to extract timing difference information from the data. The RMS time was calculated for a given event as the square root of the sum of the squared differences of every

hit time from the average hit time per event, averaged over the number of hits for that event:

$$t_{RMS}(x) = \sqrt{\frac{\sum_{i=1}^{N(x)} (t_i(x) - t_{\mu(x)})^2}{N(x)}}, \quad (3)$$

where i is an individual hit within the event x , $N(x)$ is the number of hits in event x , t_i is the recorded time of hit i , and $t_{\mu(x)}$ is the average hit time for the event.

The RMS time of flight, shown in Figure 2I, has greater resistance to dark noise fluctuations (random hits before or after an event) and was found to show greater discrimination between signal and background compared to the overall time of flight.

5.3. Distance to wall

The distribution of event vertex distance to the IWCD cylindrical tank wall, inspired by Irvine (2014), was also explored as a potential discriminating feature between neutron capture and background events. For an underground WC detector such as Super-Kamiokande, which is located approximately one kilometer underground, a greater number of background events may originate at positions nearer the detector walls due to radiation from the surrounding rock. In the simulated IWCD data, there is a slightly greater occurrence of neutron capture events in the region of 50–300 cm from the tank wall, as seen in Figure 2H.

5.4. Mean opening angle

The Cherenkov emission from relativistic photons in water is emitted on a cone with respect to the origin of radiation. The angle of emission is dependent on the kinematic properties of the incident charged particles. The mean opening angle from the event vertex varies on average for different types of particle interactions, making this metric another possible discriminant to improve neutron tagging performance. Following the definition in Irvine (2014), this mean opening angle is calculated as the mean value of the angles between every hit PMT vector and the true vertex position within the given event:

$$\Theta_{\mu}(x) = \frac{\sum_{i=1}^{N(x)} \Theta(\vec{p}_i, \vec{p}_0)}{N(x)} = \frac{\sum_{i=1}^{N(x)} \arccos\left(\frac{\vec{p}_i \cdot \vec{p}_0}{|\vec{p}_i| \cdot |\vec{p}_0|}\right)}{N(x)}, \quad (4)$$

where $\Theta_{\mu}(x)$ is the mean opening angle for the event x , $N(x)$ is the number of hits, \vec{p}_i is the (x, y, z) position of the i -th hit, \vec{p}_0 is the (x, y, z) position of the event vertex and $\Theta(\vec{p}_i, \vec{p}_0)$ is the angle between \vec{p}_i and \vec{p}_0 , computed as the quotient of the dot product by the product of their magnitudes.

The mean opening angle metric is largely influenced by the event energy. Discrimination is observed between the distributions seen in Figure 2J due to a combination between the event energy and topological distribution of the hits throughout the event. The electron events in the background dataset have lower energies, but the hits are more sparsely distributed. Evidently, the net effect is that the neutrons end up with a higher peak mean opening angle than the background events in the dataset, on average.

5.5. Consecutive hit angular RMS

Another potential neutron tagging feature discriminant, again inspired by Irvine (2014), is the root-mean-squared consecutive angle of an event. True background hits, for example from radioactive background sources, often contain spatially compact clusters of hits. On the other hand, Cherenkov photons from neutron capture events would be expected to propagate more uniformly within the average opening angle of the radiation emission cone. The RMS difference of angle between temporally consecutive hits can extract information on these angular differences between event types. The RMS angle is calculated by first sorting all PMT hits chronologically within a given event, then computing the sum of the squared differences of the angles between consecutive events from the mean consecutive angular difference, averaged over the number of hits for the event and square rooted, as

$$\Theta_{RMS}(x) = \sqrt{\frac{\sum_{i=1}^{N(x)-1} (\Theta(\vec{p}_i, \vec{p}_{i+1}) - \Theta_{\mu})^2}{N(x)}} = \sqrt{\frac{\sum_{i=1}^{N(x)-1} (\arccos\left(\frac{\vec{p}_i \cdot \vec{p}_{i+1}}{|\vec{p}_i| \cdot |\vec{p}_{i+1}|}\right) - \Theta_{\mu})^2}{N(x)}}, \quad (5)$$

where $\Theta_{RMS}(x)$ is the RMS consecutive angle for the event x , $N(x)$ is the number of hits, \vec{p}_i is the (x, y, z) position of the i -th hit, \vec{p}_{i+1} is the (x, y, z) position of next consecutive hit in time order $i+1$, Θ_{μ} is the average angle between consecutive hits in the event and $\Theta(\vec{p}_i, \vec{p}_{i+1})$ is the angle between \vec{p}_i and \vec{p}_{i+1} .

For events with more scattering, clustering and reflections, the distributions of RMS consecutive angles will be higher on average, and vice versa. Figure 2K shows that there is little difference between neutron and background signals, which is expected since our simulation uses a uniform distribution of background events. For a background source more inclusive of clustering, the discrimination extent is expected to be greater for the RMS angular metric.

5.6. Consecutive hit distance

In studying the event displays of the neutron capture and background events, it was observed that the positional distributions of hits tended to be more widespread in neutron capture events. Given two events of different type with similar numbers of hits, the neutron capture event could be reasonably well-differentiated by eye by selecting the event with greater average distance between hits. To compute the average consecutive hit distance, the hits within a given event were first sorted chronologically in time, then the Euclidean distances between consecutive hits were summed and averaged over the number of hits within the event as

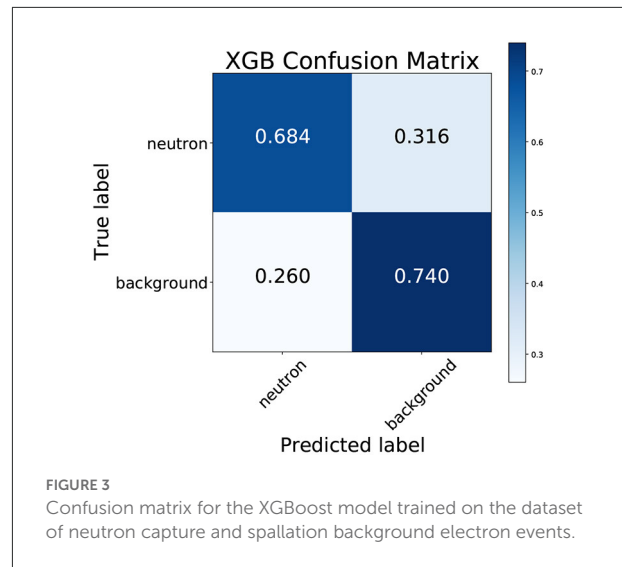
$$hd_{\mu}(x) = \frac{\sum_{i=1}^{N(x)-1} dist(\vec{p}_i, \vec{p}_{i+1})}{N(x)} = \frac{\sum_{i=1}^{N(x)-1} \sqrt{(p_{x(i)} - p_{x(i+1)})^2 + (p_{y(i)} - p_{y(i+1)})^2 + (p_{z(i)} - p_{z(i+1)})^2}}{N(x)}, \quad (6)$$

where $hd_{\mu}(x)$ is the average consecutive hit distance for the event x , $N(x)$ is the number of hits, \vec{p}_i is the (x, y, z) position of the i -th hit, \vec{p}_{i+1} is the (x, y, z) position of the next consecutive hit in time order $i+1$ and $dist(\vec{p}_i, \vec{p}_{i+1})$ is the Euclidean distance between consecutive hits.

The difference of consecutive hit distance was a good discriminator, as seen in Figure 2L. This difference in consecutive hit distance is likely due to the differing nature of the particle interactions, in which the cascade of gammas from the neutron capture leads to greater spatial separation of hits throughout the detector, on average, when compared to the electron background hits.

5.7. XGBoost results

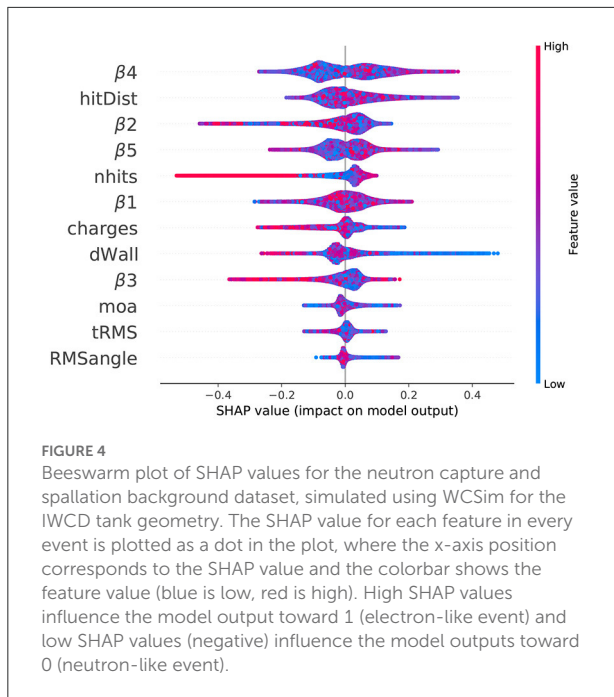
The XGBoost gradient boosting decision tree model was applied to the task of learning from the features engineered in Section 5. A grid search was applied to tune the model hyperparameters, including the maximum tree depth max_depth , the minimum tree node weight min_child_weight , the training data subsampling ratio $subsample$, the tree column subsampling ratio $colsample_bytree$, and the learning rate eta . The grid search sequentially iterated over relating parameters pairs and applied four-fold cross-validation to improve outcome reliability. The relating pairs were max_depth and min_child_weight , followed by $subsample$ and $colsample_bytree$. The learning rate was adjusted independently. For each hyperparameter combination, XGBoost’s native cross-validation function was used to train the model over a maximum of 1,250 boosting rounds, and early stopping was used to cancel model training if performance did not improve over twenty consecutive rounds.



Applying this technique, the optimal tree complexity was found with max_depth of 11 and a min_child_weight of 1, the optimal sampling ratios were found with a $subsample$ ratio of 0.7 and a $colsample_bytree$ ratio of 1.0, and the learning rate was tuned to 0.007. The optimized XGBoost model was then trained on a consistent 80% training dataset, optimized against an independent 10% validation dataset and tested against a 10% holdout test set. The model obtained train, validation and test accuracies of 73.0, 71.5, and 71.4%, respectively, and an ROC AUC score of 0.784. Although the training accuracies are generally slightly higher than the test accuracy, the extent of overfitting was not too severe and may be decreased by using a smaller number for early stopping. The XGBoost model construction for any of the 80% training sets was found to take ~45–60 min, depending on the number of trees constructed before early stopping. Figure 3 displays the confusion matrix, which shows that the true positive rate (neutron sensitivity) is significantly lower than the true negative rate (neutron specificity).

SHAP was used to understand the relative importances of the dozen features contributing to the XGBoost model. The SHAP values are applicable both locally, to a single event, and globally, to a conglomerate of events. While various visualizations using SHAP are possible, the beeswarm plot, in particular, is useful in showing the range and density of SHAP values for individual features.

In the beeswarm plot, it is hard to see the distribution as it has a density of points as a color for the value. The main reason for introducing it here is to see the reveals a notable difference between the lower order ($\beta_1, \beta_2, \beta_3$) and higher order (β_4, β_5) isotropy parameters. Figure 4 shows the beeswarm plot over all events in the neutron capture and spallation electron background dataset. For this plot, the SHAP value for each feature in every event is plotted as a single dot. Bulges in a row indicate areas of larger density. Higher SHAP values



influence the model output toward 1 (electron-like event) and low SHAP values (negative) influence the model outputs toward 0 (neutron-like event). The features are arranged on the vertical axis by feature importance, with the most important features (by average absolute SHAP value) on the top and the least important features on the bottom. Each feature value is plotted with a color corresponding to its position within its numeric range.

Several distinctive patterns from Figure 4 are discernible. For a high number of hits, the SHAP value is uniformly negative. Correspondingly, in Figure 2A, it is clear that events with more than approximately 100 hits are uniformly neutron events (top-left plot). For the wall distance, it is clear from Figure 2H that there are a slight over-representation of background events at distances close to the wall. The XGBoost model clearly notices this difference, as events with lower wall distances mostly have higher SHAP values, meaning the model output value is pushed higher to 1 (electron-like event).

The beeswarm Figure 4 also reveals a notable difference between the lower order (β_1 , β_2 , β_3) and higher-order (β_4 , β_5) isotropy parameters. β_1 , β_2 , and β_3 both have single mode representations in the beeswarm plot, in which there is a single bulge. Higher values for these parameters also attribute the output toward a neutron classification, on average. This correspondence may be seen by the feature differences of Figures 2C–G. Alternately, β_4 and β_5 have two main modes (bulges) in the SHAP value beeswarm plot, indicating two main regions with SHAP values of a similar range. For β_5 , a clear distinction is seen between lower values of β_5 , attributed toward neutron events, and higher values of β_5 , attributed toward electron events. While this difference is clear, the SHAP values

themselves are lower, showing a smaller output impact. This small difference is observable in Figure 2G for β_5 .

The β_4 parameter has a similar double-modal pattern in the beeswarm plot, but the attributed difference is smaller for lower and higher values of the parameter. However, β_4 still has the greatest average absolute SHAP value, and therefore the greatest average impact on the model output. In general, β_4 , mean consecutive hit distance, β_2 , β_5 , and number of hits, respectively were the top five most important features in determining event outcomes.

6. Graph neural network application

In this study, the PyTorch Geometric (PyG) library was used to apply graph neural network models to the IWCD dataset (Fey and Lenssen, 2019). This particular library was chosen for its ease of use, breadth of graph network models available, data loading tools and GPU support. During training, at regular intervals, the model was applied to the validation dataset to check for under fitting or overfitting. After the model was trained, it was applied on the test dataset and evaluation metrics were computed. Model parameters were updated using Adam optimization (Kingma and Ba, 2014) with cross-entropy loss. Training was carried out on a Quadro P2000 GPU.

6.1. Graph convolutional network (GCN)

For a neutron capture or background event, the hit PMTs may be represented as graph nodes, with each node containing the features of hit time, deposited charge and the three-dimensional position and orientation of the hit PMT. Since the number of hits varies for every event, the graphs could either vary in size (non-padded graph) or zero padding could be added. Graph padding, along with edge weighting and node connectivity, were three hyperparameters of graph construction investigated in this research. Within the GCN framework, model performance was compared against padded vs. non-padded graphs, edge weighted (inversely proportional to distance) vs. uniform weights, and the fully connected vs. k nearest neighbor graph.

To begin, the GCN model was tested on graphs constructed using a padded, fully connected (every node connected to every other node) representation with all edge weightings set to a value of one. This setting was used to adjust parameters of the GCN architecture, leading to the configuration of two alternating layers of GCN convolutional filtering and activation computation, with 24 and 8 compute nodes in the first and second hidden layers, respectively. This was followed by max pooling and the log softmax output from a

fully connected layer with two neurons in the output layer. GCN results were obtained by training with a batch size of 32, learning rate of 0.0003 and learning rate decay of 0.001%.

The first graph construction comparison tested whether the GCN model learned better on padded graphs or variable-size graphs without padding. The results of training the GCN model for *padded* and *non-padded*, fully connected graphs with uniform edge weightings are shown in the first two rows of [Table 1](#). The performance was higher with for padded graphs, with an average test accuracy improvement of 3.2%.

While accuracy was higher for the padded graphs, the runtime was also considerably longer due to the significantly greater number of connections and message passing operations in the padded graphs. Run times were recorded per epoch, where an epoch is one entire transit of the training data through the algorithm. Per epoch, the padded graphs took 6 h to train, while the non-padded graphs took only 14 min. The non-padded graph GCN model was trained over 75 epochs and ~ 17 h, while the padded GCN model was trained over 5 epochs and ~ 30 h. A higher number of epochs was not found to improve the performance for either model. A summary of runtimes is presented in [Table 1](#).

Next, edge weighting was tested for the GCN model to see if edge values related to physical distance could provide a learning advantage over uniform edge weightings set to a tensor of ones. The results of training the GCN on fully connected, padded, inverse-distance edge weighted graphs is shown in the last row of [Table 1](#). Runtimes are nearly identical to the same model with fixed edge weights. With distance weighted edges, the test accuracy was 1.7% lower than the corresponding result for graphs with uniform edge weights. The edge weightings possibly overcomplicate the GCN model on the scale of the $\sim 10^5$ node connections for a given event.

Overall, the GCN model was found to perform best on static, fully connected, uniform edge weighted graphs. The results are shown in [Table 1](#). This GCN configuration has comparable metrics to the highest likelihood baseline, with 0.6% higher accuracy on the spallation set. Moreover, training results were nearly identical whether the model was fed only the hits (and charges) data, or if the position and orientation data was included along with the hits and charges. Therefore, it appears that the GCN model was largely learning to classify events based on the trivial number of hits, and that it failed to significantly learn from the geometric differences of neutron capture to electron background hit patterns. Adding any additional network layers to the GCN models was also found to worsen performance, presumably as the extra filtering step oversmooths the node representations.

6.2. Dynamic graph convolutional neural network

The DGCNN model was the next graph network model applied to the particle classification task. Described in [Section 3.3.2](#), the DGCNN model was selected for its ability to learn from point cloud data specifically. The network architecture configuration was set to the default from the PyTorch Geometric example documentation, which consisted of the following: two dynamic edge convolution layers followed by a fully connected layer, a global max pooling layer and a final MLP to yield the output class probabilities. The first edge convolution layer applied an MLP on input node features with three layers of 64 compute units each. The second edge convolution layer took the output of the first as input and applied an MLP with 128 activation units. In both cases, the MLP was applied to every node pair ($n * 2$ pairs for n nodes in an event) over the k -nn (k nearest neighbor) graph representation of each node, and the representations were updated by pooling the learned edge features. After the EdgeConv blocks, a fully connected layer concatenated the 64 and 128 unit features from the first two dynamic edge convolution layers and yielded 1,024 activations. Global max pooling was applied over the n nodes to reduce the representation from $n * 1,024$ to only 1,024. The final MLP then passed this information into final layers of 512, 256, and 2 activation nodes, respectively and the softmax of the output was applied to calculate the output the binary classification probabilities. Note that for every fully connected layer throughout the network, the activations were calculated using the ReLU activation function and batch normalization ([Ioffe and Szegedy, 2015](#)) to reduce overfitting. The model description is represented by [Figure 5](#).

With the fixed architecture as described above, the number of nearest neighbors k in the DGCNN dynamic edge convolution blocks were adjusted over multiple runs to compare performance. [Table 2](#) shows the results of applying the DGCNN model on the spallation background dataset with the k hyperparameter varying from 10 to 30 in increments of 5. The resulting accuracies were largely the same for $k = 15$ to $k = 30$, while $k = 10$ neighbors led to a slightly lower accuracy. Among the range of $k = 15$ to $k = 30$, $k = 25$ yielded the highest ROC AUC score of 0.797, although the 0.001 difference compared to the other k values in the range was not necessarily statistically significant.

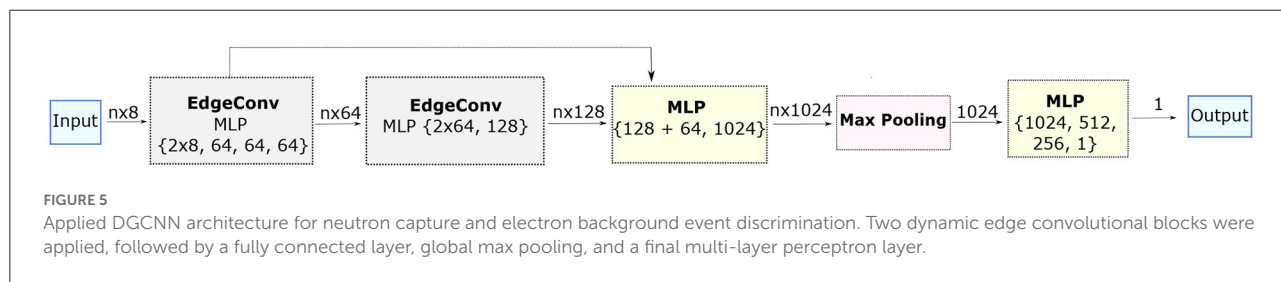
Regarding the statistical sensitivity, for the 80,000 validation events in each sample the statistical uncertainty is 0.4%. Since the same datasets are being reused in changing the k values, there is some cancelation of uncertainty due to correlation through using the exact same data. A conservative estimate of 0.4% uncertainty could be used here, since relying on correlation by

TABLE 1 GCN model applied to padded and non-padded, fully connected, uniformly edge weighted, and inverse square distance ($1/d^2$) weighted graphs for the simulated IWCD neutron capture and background datasets.

Type	Train accuracy	Validation accuracy	Test accuracy	ROC AUC
Padded	61.3	63.1	63.1	0.667
Non-padded	58.5	59.8	59.9	0.628
Padded ($1/d^2$)	59.7	61.3	61.4	0.632

TABLE 2 DGCNN model classification accuracies for variations of the number of nearest neighbors k in the DGCNN dynamic edge convolution blocks from 10 to 30 in increments of 5.

k neighbors	Train accuracy	Validation accuracy	Test accuracy	ROC AUC
10	70.9	72.0	71.9	0.792
15	71.8	72.2	72.3	0.796
20	71.7	72.3	72.3	0.796
25	71.8	72.4	72.4	0.797
30	71.4	72.4	72.4	0.796



not having generated enough data to compare the methods using independent data may mean fitting to some peculiarity of the dataset, rather than a generally useful difference.

While there was minimal performance difference for the range of $k = 15$ to $k = 30$, there was however a difference in the training times, as shown in Table 3. This was expected as, for $n f$ -dimensional input nodes, an $n * k * a_n$ -dimensional tensor is generated before pooling across the neighboring edge features for every dynamic edge convolution block. Therefore, the total number of training parameters increases significantly for every increment of nearest neighbors k . There was a sharp increase in training time after about $k = 20$ and more than a doubling in overall training time from $k = 10$ to $k = 30$.

Given the results in Table 2 and the processing times required, $k = 15$ is a good compromise between training time and classification accuracy. However, when training time is not a significant impediment, $k = 25$ might be used to optimize results.

7. Discussion

For all models, consistent training, validation, and test datasets were constructed in an 80, 10, and 10% ratio. Models

TABLE 3 Comparison of training times for the different models applied in this study, sorted in ascending order by training time per epoch.

Model	Epochs	Total runtime (min)	Time per epoch (min)
XGBoost	1,450	50	0.036
DGCNN ($k = 10$)	25	1,980	1.32
DGCNN ($k = 15$)	25	2,100	1.4
DGCNN ($k = 20$)	25	2,700	1.8
DGCNN ($k = 25$)	25	3,420	2.28
DGCNN ($k = 30$)	25	3,960	2.64
GCN (non-padded)	75	1,020	13.6
Likelihood Ratio	1	80	80
GCN (padded)	5	1,800	360

were optimized against the validation data and metrics were reported for the holdout test dataset, ensuring that differences in model performance were not due to random distributions of the data. Compared to the likelihood statistical baseline, the DGCNN model results in Table 2 showed an accuracy

TABLE 4 Overall accuracies for neutron capture vs. electron background classification for the likelihood analysis (Likelihood), XGBoost, GCN, and DGCNN methods.

Dataset background source	Likelihood	XGBoost	GCN	DGCNN
Spallation	62.5	71.4	63.1	72.4

improvement of 9.9%. The $\sim 10\%$ classification accuracy improvement strongly indicates the capability of the DGCNN model to learn from event topology and other, more subtle factors than the number of hits and overall sum of charges within the event. The dynamic method of graph construction with the DGCNN model, which shuffles the groupings of every node with its other nearest neighbor nodes in semantic space, allows the diffusion of nonlocal information throughout the graph. This ostensibly allows the DGCNN model to learn global event topology in a way which the GCN model, restricted to operating over fixed input graphs, was not able to.

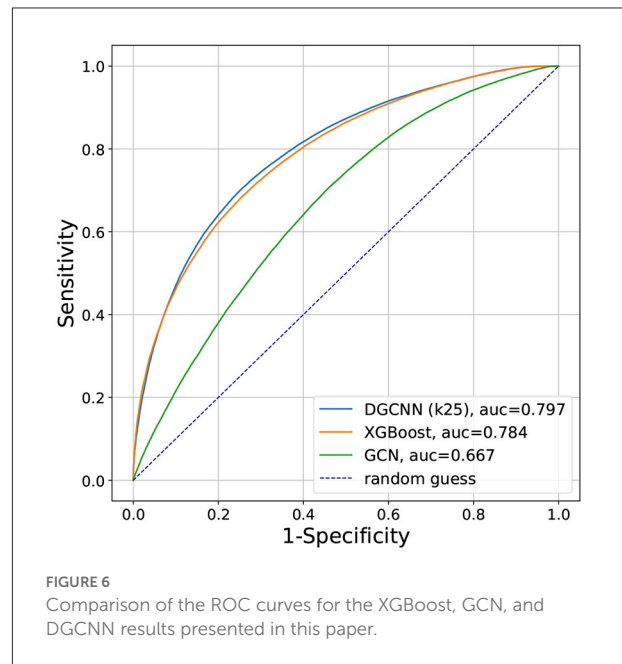
Overall, the DGCNN also slightly outperformed the best XGBoost model, representing an improvement in accuracy of 0.7% and ROC AUC score of 0.007. The test accuracy results for all approaches undertaken in this study, including the likelihood baseline analysis, XGBoost with feature engineering and the GCN and DGCNN models are presented in Table 4. The best accuracy for neutron vs. background separation was 72.4% using DGCNN.

The receiver operating characteristic (ROC) curves, which plots the true positive rate (sensitivity) against the false positive rate (1—specificity) for a binary classification problem, are shown in Figure 6 for the different machine learning methods studied in this paper. The ROC AUC (area under the curve) from XGBoost is 0.784, from GCN is 0.667, and from DGCNN with $k = 25$ is 0.797, showing that consistent with the accuracies presented earlier, the DGCNN had the best performance.

8. Conclusions

This paper has presented a search to improve the classification performance of neutron capture vs. background identification in WC detectors using techniques in machine learning. To provide a performance baseline, a statistical model was applied to classify events using maximum likelihood of kernel density estimates of the main event type discriminants, namely the number of hits and charge sums. The baseline accuracy was found to be 62.4%.

Next, a series of features were engineered from the datasets. Besides number of hits and charge sums, the beta parameters β_1 – β_5 were created to capture event isotropy. The mean opening angle, event vertex distance to wall, RMS consecutive hit angle and mean consecutive hit distance were also computed to summarize event topology and the RMS event time was added to capture timing discrimination. Gradient boosted decision



trees were applied on these engineered features using the XGBoost algorithm. The XGBoost model hyperparameters were tuned using grid search, yielding a test accuracy of 71.4% which represented an 8.9% improvement over the baseline approach. SHAP analysis of these model outputs revealed useful information. The β_2 , β_4 , β_5 , number of hits and consecutive hit distance parameters were consistently rated most important, as measured by the mean absolute SHAP value.

Drawbacks to the XGBoost and the feature engineering approach include preprocessing time to calculate the feature values and the fact that the calculation of several features relies on the event vertex position. For this research, the true vertex position was taken from the simulation information, but in reality a vertex reconstruction algorithm would need to be used, introducing some uncertainties into the equations. Future studies could endeavor to reduce the bias in the engineered features by smearing the true vertex positions over a range of values.

As an alternative approach, deep learning was implemented via the GCN and DGCNN graph neural network models. The GCN model was tested with a variety of graph construction approaches, including static vs. non-static graphs, uniform edge weighting vs. scaled edge weights, and fully connected vs. partially connected graphs. Of all these cases, the best

test accuracy obtained was 63.1% using the fully connected, zero padded, uniform edge weighted graphs, which was nearly identical to the baseline likelihood accuracy.

The DGCNN model, however, was found to have significantly improved neutron tagging performance above the baseline accuracy. The DGCNN number of neighbors hyperparameter k was tuned, and the reported accuracy was found to be 72.4%, representing an improvement of 9.9% over the likelihood analysis. Thus, DGCNN slightly outperformed XGBoost on the classification of neutron vs. background. DGCNN also retains the advantage of not requiring any preprocessing or prior knowledge. On the other hand, XGBoost provides a much greater level of model interpretability. Furthermore, once the engineered features have been computed, the training time of XGBoost for the datasets used in this study was within the range of only 45 min to 1 h, much faster than the DGCNN model which took from 30 to over 60 h, depending on the value of k . However, DGCNN was trained over only a single GPU, and using multiple GPUs could reduce the runtime significantly. Table 4 shows the overall results of XGBoost, GCN, and DGCNN compared to the likelihood baseline.

Overall, both XGBoost with feature engineering and DGCNN show promise in improving neutron tagging efficiency in WC detectors. In particular, the application of these methods in the IWCD might help reduce systematic uncertainties for the Hyper-Kamiokande detector, which it turn could advance our understanding of neutrino physics and the Standard Model itself. In future, the network architecture of the DGCNN model could be further optimized.

For practical purposes, given that these models were developed for data simulation, another reasonable next step would include the deployment of these models for neutron tagging in active WC detectors. This would test if the models are transferable for real use cases. Also, these models could be incorporated into a pipeline that tests for the coincidence of neutron capture and positron rings within a timescale indicative of neutrino inverse beta decay. While the development of improved neutron tagging is desirable, the ultimate goal is to trace back to the originating neutrino to probe deeper into the unknowns of neutrino physics. An end-to-end network could thus be deployed using the neutron tagging models developed in this research to better identify the neutrinos themselves in the overarching process of the neutrino inverse beta decay.

Data availability statement

The original contributions presented in the study are included in the article/supplementary materials, further

inquiries can be directed to the corresponding author (bl.jamieson@uwinnipeg.ca).

Author contributions

This paper presents the research conducted by MS, whose thesis this paper is based on Stubbs (2021). BJ, SR, JW, NP, RA, PP, and WF contributed to the development of the research at weekly meetings. The initial implementation of the GNN code was prepared by JW, and datasets were prepared by NP. All authors contributed to the article and approved the submitted version.

Funding

The funding for this research is from the Canadian National Science and Engineering Council (NSERC). Production of the simulation datasets was done with the support of Compute Canada resources.

Acknowledgments

We acknowledge the WatChMaL, Super-Kamiokande, and Hyper-Kamiokande collaborations, on whose shoulders we stand in coming up with the idea for this study, and with whom many of the authors are collaborators. This research was enabled in part by support provided by Cedar Compute Cluster (<https://docs.alliancecan.ca/wiki/Cedar>) and the Digital Research Alliance of Canada (alliancecan.ca).

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.

References

- Agostinelli, S. (2003). Geant4-a simulation toolkit. *Nucl. Instrum. Method. Phys. Res. Sec. A* 506, 250–303. doi: 10.1016/S0168-9002(03)01368-8
- Andrews, M., Paulini, M., Gleyzer, S., and Poczoz, B. (2020). End-to-end physics event classification with CMS open data: applying image-based deep learning to detector data for the direct classification of collision events at the LHC. *Comput. Softw. Big Sci.* 4:6. doi: 10.1007/s41781-020-00038-8
- Ankowski, A. M., Benhar, O., Mori, T., Yamaguchi, R., and Sakuda, M. (2012). Analysis of γ -ray production in neutral-current neutrino-oxygen interactions at energies above 200 MeV. *Phys. Rev. Lett.* 108:052505. doi: 10.1103/PhysRevLett.108.052505
- ATLAS Collaboration (2017). *Identification of Jets Containing b-Hadrons With Recurrent Neural Networks at the ATLAS Experiment*. Technical report, CERN, Geneva. Available online at: <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2017-003> (accessed January 26, 2022).
- ATLAS Collaboration (2019). *Convolutional Neural Networks With Event Images for Pileup Mitigation with the ATLAS Detector*. Technical report, CERN, Geneva. Available online at: <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2019-028> (accessed January 26, 2022).
- Beacom, J. F., and Vagins, M. R. (2004). Antineutrino spectroscopy with large water Cherenkov detectors. *Phys. Rev. Lett.* 93:171101. doi: 10.1103/PhysRevLett.93.171101
- Bernard, L. (2019). “Spallation background in the Super-Kamiokande experiment,” in *Super-Kamiokande Collaboration ICHEP Conference, Neutrino Session*.
- Bhattacharya, S., Maddikunta, P. K. R., Kaluri, R., Singh, S., Gadekallu, T. R., Alazab, M., et al. (2020). A novel PCA-firefly based XGBoost classification model for intrusion detection in networks using GPU. *Electronics* 9:219. doi: 10.3390/electronics9020219
- Bourilkov, D. (2019). Machine and deep learning applications in particle physics. *Int. J. Modern Phys. A* 34:1930019. doi: 10.1142/S0217751X19300199
- Brun, R., and Rademakers, F. (1997). R.O.O.T.-an object oriented data analysis framework. *NIM A* 389, 81–86. doi: 10.1016/S0168-9002(97)00048-X
- Bruna, J., Zaremba, W., Szlam, A., and LeCun, Y. (2013). Spectral networks and locally connected networks on graphs. *arXiv [Preprint]*.arXiv: 1312.6203. doi: 10.48550/ARXIV.1312.6203
- Carleo, G., Cirac, I., Cranmer, K., Daudet, L., Schuld, M., Tishby, N., et al. (2019). Machine learning and the physical sciences. *Rev. Mod. Phys.* 91:045002. doi: 10.1103/RevModPhys.91.045002
- Chen, T., and Guestrin, C. (2016). “XGBoost: a scalable tree boosting system,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16* (New York, NY: Association for Computing Machinery), 785–794. doi: 10.1145/2939672.2939785
- Chen, T., and He, T. (2015). “Higgs boson discovery with boosted trees,” in *Proceedings of the NIPS 2014 Workshop on High-Energy Physics and Machine Learning, Vol. 42*, eds G. Cowan, C. Germain, I. Guyon, B. Kégl, and D. Rousseau (Montreal, QC: PMLR), 69–80.
- Choma, N., Monti, F., Gerhardt, L., Palczewski, T., Ronagli, Z., Prabhat, P., et al. (2018). “Graph neural networks for icecube signal classification,” in *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*. (Orlando, FL), 386–391. doi: 10.1109/ICMLA.2018.00064
- Cornell, A. S., Doorsamy, W., Fuks, B., Harmsen, G., and Mason, L. (2022). Boosted decision trees in the era of new physics: a smuon analysis case study. *J. High Ener. Phys.* 2022, 15. doi: 10.1007/JHEP04(2022)015
- Dunmore, J. A. (2004). *The separation of CC and NC events in the sudbury neutrino observatory* (Ph.D. thesis). Oxford.
- Fernández, P. (2016). Status of GADZOOKS!: neutron tagging in super-Kamiokande. *Nucl. Part. Phys. Proc.* 273–275, 353–360. doi: 10.1016/j.nuclphysbps.2015.09.050
- Fey, M., and Lenssen, J. E. (2019). Fast graph representation learning with PyTorch geometric. *arXiv [Preprint]*.arXiv: 1903.02428. doi: 10.48550/ARXIV.1903.02428
- Gligorov, V. V., and Williams, M. (2013). Efficient, reliable and fast high-level triggering using a bonsai boosted decision tree. *J. Instrumen.* 8:P02013. doi: 10.1088/1748-0221/8/02/P02013
- Guest, D., Cranmer, K., and Whiteson, D. (2018). Deep learning and its application to LHC physics. *Annu. Rev. Nucl. Part. Sci.* 68, 161–181. doi: 10.1146/annurev-nucl-101917-021019
- Hammond, D. K., Vandergheynst, P., and Gribonval, R. (2011). Wavelets on graphs via spectral graph theory. *Appl. Comput. Harm. Anal.* 30, 129–150. doi: 10.1016/j.acha.2010.04.005
- Ioffe, S., and Szegedy, C. (2015). “Batch normalization: accelerating deep network training by reducing internal covariate shift,” in *International Conference on Machine Learning* (Lille), 448–456.
- Irvine, T. J. (2014). *Development of neutron-tagging techniques and application to atmospheric neutrino oscillation analysis in Super-Kamiokande* (Ph.D. thesis). University of Tokyo, Tokyo, Japan.
- Kingma, D. P., and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv [Preprint]*.arXiv: 1412.6980. doi: 10.48550/ARXIV.1412.6980
- Kipf, T., and Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *Arxiv: abs/1609.02907*.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2017). *ImageNet Classification with Deep Convolutional Neural Networks*. New York, NY: Association for Computing Machinery. doi: 10.1145/3065386
- Li, S. W., and Beacom, J. F. (2014). First calculation of cosmic-ray MUON spallation backgrounds for MEV astrophysical neutrino signals in super-Kamiokande. *Phys. Rev. C* 89:045801. doi: 10.1103/PhysRevC.89.045801
- Lundberg, S. M., and Lee, S.-I. (2017). “A unified approach to interpreting model predictions,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, (Red Hook, NY: Curran Associates Inc.), 4768–4777.
- Macaluso, S. and Shih, D. (2018). Pulling out all the tops with computer vision and deep learning. *J. High Ener. Phys.* 2018:121. doi: 10.1007/JHEP10(2018)121
- Michael, D., Bresson, X., and Vandegheynst, P. (2016). Convolutional neural networks on graphs with fast localized spectral filtering. *arXiv [Preprint]*.arXiv: 1606.09375. doi: 10.48550/ARXIV.1606.09375
- Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., and Khudanpur, S. (2010). “Recurrent neural network based language model,” in *Interspeech* (Makuhari), 2, 1045–1048. doi: 10.21437/Interspeech.2010-343
- Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., and Bronstein, M. M. (2017). “Geometric deep learning on graphs and manifolds using mixture model CNNs,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5115–5124. doi: 10.1109/CVPR.2017.576
- O’Sullivan, E. (2021). Water Cherenkov detector simulation (WCSSIM). Available online at: <https://github.com/WCSim/WCSim>
- Parsa, A. B., Movahedi, A., Taghipour, H., Derrible, S., and Mohammadian, A. K. (2020). Toward safer highways, application of XGBoost and SHAP for real-time accident detection and feature analysis. *Accid. Anal. Prevent.* 136:105405. doi: 10.1016/j.aap.2019.105405
- Proto-Collaboration, H.-K., Abe, K., Aihara, H., Aimi, A., Andreopoulos, A. C., et al. (2018). Hyper-kamiokande design report. *arXiv:1805.04163*. doi: 10.48550/arXiv.1805.04163
- Qu, H., and Gouskos, L. (2020). Jet tagging via particle clouds. *Phys. Rev. D* 101:056019. doi: 10.1103/PhysRevD.101.056019
- 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
- Roe, B. P., Yang, H.-J., Zhu, J., Liu, Y., Stancu, I., and McGregor, G. (2005). Boosted decision trees, an alternative to artificial neural networks. *Nucl. Instrum. Methods A* 543, 577–584. doi: 10.1016/j.nima.2004.12.018
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. (2008). The graph neural network model. *IEEE Trans. Neural Netw.* 20, 61–80. doi: 10.1109/TNN.2008.2005605
- Shapley, L. S. (1953). Stochastic games. *Proc. Natl. Acad. Sci. U.S.A.* 39, 1095–1100. doi: 10.1073/pnas.39.10.1953
- Shlomi, J., Battaglia, P., and Vlimant, J.-R. (2021). Graph neural networks in particle physics. *Mach. Learn.* 2:021001. doi: 10.1088/2632-2153/abbf9a
- Sirunyan, A. M., Tumasyan, A., Adam, W., Ambrogio, F., Bergauer, T., Dragicevic, M., et al. (2020). Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques. *J. Instrument.* 15, P06005. doi: 10.1088/1748-0221/15/06/p06005
- Stubbs, M. (2021). *Using machine learning to improve neutron tagging efficiency in water Cherenkov detectors* (Master’s thesis). University of Winnipeg, Winnipeg, MB, Canada.

- Tahmassebi, A., Wengert, G. J., Helbich, T. H., Bago-Horvath, Z., Alaei, S., Bartsch, R., et al. (2019). Impact of machine learning with multiparametric magnetic resonance imaging of the breast for early prediction of response to neoadjuvant chemotherapy and survival outcomes in breast cancer patients. *Invest. Radiol.* 54, 110–117. doi: 10.1097/RLI.0000000000000518
- Wang, Y., Sun, Y., Liu, Z., Sarma, S. E., Bronstein, M. M., and Solomon, J. M. (2019). Dynamic graph CNN for learning on point clouds. *ACM Trans. Graph.* 38, 1–12. doi: 10.1145/3326362
- Watanabe, H., Zhang, H., Abe, K., Hayato, Y., Iida, T., Ikeda, M., et al. (2009). First study of neutron tagging with a water Cherenkov detector. *Astropart. Phys.* 31, 320–328. doi: 10.1016/j.astropartphys.2009.03.002
- Wilson, J. R. (2015). An experimental review of solar neutrinos. eds. R. Aumann and S. Hart, North Holland. *arXiv [Preprint]*. arXiv: 1504.04281. doi: 10.48550/ARXIV.1504.04281
- Winter, E. (2002). “The shapley value,” in *Handbook of Game Theory with Economic Applications*, Vol. 3, eds R. Aumann and S. Hart (Amsterdam: Elsevier), 2025–2054. doi: 10.1016/S1574-0005(02)03016-3
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., et al. (2020). Graph neural networks: a review of methods and applications. *AI Open* 1, 57–81. doi: 10.1016/j.aiopen.2021.01.001