

LArTPC hit-based topology classification with quantum machine learning and symmetry considerations

Callum Duffy¹, Marcin Jastrzebski¹, Stefano Vergani^{1,*}, Leigh H. Whitehead², Ryan Cross³, Andrew Blake⁴, Sarah Malik¹ and John Marshall³

¹*Department of Physics and Astronomy, University College London, London, United Kingdom*

²*Cavendish Laboratory, University of Cambridge, Cambridge, United Kingdom*

³*Department of Physics, University of Warwick, Coventry, United Kingdom*

⁴*Physics Department, Lancaster University, Lancaster, United Kingdom*



(Received 22 May 2025; accepted 2 September 2025; published 12 November 2025)

We present a new approach to separate tracklike and showerlike topologies in liquid argon time projection chamber (LArTPC) experiments for neutrino physics using quantum machine learning. Effective reconstruction of neutrino events in LArTPCs requires accurate and granular information about the energy deposited in the detector. These energy deposits can be viewed as 2D images. Simulated data from the MicroBooNE experiment and a simple custom dataset are used to perform pixel-level classification of the underlying particle topology. Images of the events have been studied by creating small patches around each pixel to characterize its topology based on its immediate neighborhood. This classification is achieved using convolution-based learning models, including quantum-enhanced architectures known as quanvolutional neural networks. The quanvolutional networks are extended to symmetries beyond translation. Rotational symmetry has been incorporated into a subset of the models. This study demonstrates that quantum-enhanced models perform better than their classical counterparts with a comparable number of parameters, but are outperformed by classical models with two orders of magnitude more parameters. The inclusion of rotation symmetry appears beneficial only in a small number of cases and remains to be explored further. Possible future use of quantum machine learning in the reconstruction phase is discussed, with emphasis on future LArTPC experiments such as Deep Underground Neutrino Experiment (DUNE)-far detector (FD).

DOI: [10.1103/byy5-zk73](https://doi.org/10.1103/byy5-zk73)

I. INTRODUCTION

Liquid argon time projection chamber (LArTPC) detectors are commonly used in neutrino physics and represent the current state of the art of in the field. Although there are differences in the design, they normally comprise a cryostat filled with high-purity liquid argon kept at 87 K, a system to apply an internal electric field, a readout mechanism made with wires, and a photon detection system. In LArTPCs, charged particles interact with argon nuclei, liberating electrons through ionization. These electrons drift in an electric field toward readout wires, where they are detected. Combining the wire number with the time of detection, especially if the scintillation light is used as a trigger, gives spatial and time information in a 2D view.

Calorimetric information is added to each 2D view. The interactions of neutrinos with the argon nuclei are inferred by the presence of charged particles produced in the interaction emerging from a common vertex. An example of a 2D image produced by a neutrino interaction in a LArTPC experiment is shown in Fig. 1. The Micro Booster Neutrino Experiment (MicroBooNE) [1] employed an 85 metric tons LArTPC detector to perform precision physics measurements and collected five years of data from 2016 to 2021. The experiment's primary scientific goals were to solve the puzzle of the MiniBooNE low energy excess [2], to measure different neutrino cross sections, and to search for astrophysical phenomena. It was part of the wider Short Baseline Neutrino (SBN) [3], an ambitious program at Fermi National Accelerator Laboratory (FNAL) aimed at investigating eV-scale sterile neutrinos, neutrino-nucleus interactions at the GeV energy scale, and the advancement of the liquid argon detector technology. The DUNE [4] represents the next generation of neutrino physics experiments, and is partly located at the Long-Baseline Neutrino Facility (LBNF) at FNAL and partly at Sanford Underground Research Facility (SURF). It aims to answer

*Contact author: s.vergani@ucl.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. An example neutrino interaction event viewed in a LArTPC display from the MicroBooNE experiment. Six particles can be seen emerging from the interaction vertex: two tracklike (muon and proton) and four showerlike (photons). Figure adapted from Ref. [16].

fundamental questions in particle physics, such as the neutrino mass ordering [5], the value of the charge-parity (CP)-violating phase [6], the formation of black holes after a supernova explosion [7], and the hypothetical proton decay [8]. The FD, one of the components of DUNE, will also utilize LArTPC technology [9], which has successfully been used in a number of recent experiments [10–14]. Of particular interest in long-baseline neutrino oscillation experiments is the ratio of electron to muon neutrinos arriving at the detector, making the primary task in such experiments reconstructing and identifying events as muon neutrino or electron neutrino events. To do so, the spatial details of the event captured need to be understood thoroughly. An accurate classification of each of the signals received is thus paramount. This work will make use of the MicroBooNE open dataset [15], but the longer-term objective of the project is to provide tools for processing data recorded by the DUNE-FD.

Pandora [17–19], typically run as part of LArSoft [20] workflows, is one of the most popular reconstruction frameworks for LArTPC experiments. It has a modular approach where each algorithm performs a specific, and typically small, task. In this way, it is easier to guarantee that each step is done correctly. One of the most important steps is to separate track and showerlike topologies at the hit level, a crucial task to perform particle identification and event reconstruction. In order to do so, Pandora makes extensive use of machine learning (ML) techniques throughout the reconstruction chains. ML and deep learning (DL) have also been utilized in the MicroBooNE experiment [21], and there are several proposed applications of these techniques in DUNE [22–25]. Different approaches to separate tracklike particles traveling through dense electromagnetic showers have been explored [26–29], but it remains one of the many open reconstruction problems where an increase in

the performance is still sought after. The presented work contributes to this important subroutine.

The exploration of ever more capable machine learning models and recent rapid advancements in quantum technologies led to the birth of quantum machine learning (QML). One common critique of the current QML research is that models which appear successful are often benchmarked on simple problems [30,31]. This study is interested in a real open problem in experimental neutrino physics, which remains unsolved by current machine learning methods. Previous QML uses in LArTPC experiments have been proposed in the context of event classification [32,33] as well as event generation [34]. Together with many other early studies investigating the use of quantum computers in high energy physics (HEP), the community is an active contributor to the development of quantum technologies while being posed to be their great beneficiary [35].

A high-level introduction to the field of QML is given in Sec. II. In Sec. III, a detailed description of the LArTPC datasets used can be found. Section IV discusses model design concepts, and the quantum and classical architectures used are described in Sec. V. In Sec. VI, one can find the results of applying the developed models to the problems.

II. QUANTUM MACHINE LEARNING

Quantum machine learning has emerged as one of the most promising families of quantum algorithms. It has garnered significant attention due to the success and widespread adoption of classical machine learning and the compatibility of QML methods with noisy intermediate-scale quantum (NISQ) [36] devices. Combining these factors makes QML a particularly attractive avenue for exploring quantum computational advantages in practical settings.

While the search for practical quantum advantages in QML remains an active area of research, significant strides have been made theoretically [37–42]. These advances, combined with scrupulous dequantization¹ studies [43], continue to refine our understanding of the field’s potential and its limitations. Empirically, small hints of potential advantage of quantum models appear in the form of better performance or similar performance with fewer parameters [33,44–46]. It is crucial, however, to state that benchmarking quantum models against classical ones is, in general, a difficult task [31]. We find ourselves in a time where classical simulations of QML pipelines are prohibitively complex beyond 10s of qubits but quantum hardware is not yet ready to train large quantum models, either. What we can do in the meantime is understand the models we work with well and test them on challenging, real-life problems.

¹Dequantization refers to finding classical algorithms with complexity scaling matching a proposed quantum algorithm.

A critical consideration in the design of quantum learning models is their inductive bias, which refers to the assumptions embedded in the model architecture that influence its generalization capabilities. This principle is the focus of geometric [47–49] and scientific [50–52] QML, where domain knowledge is incorporated into model design to enhance performance and interpretability.

The model used in this study is the well-known quanvolutional neural network which has been extended, for the first time, to include symmetries beyond translation. Quanvolutional neural networks allow one to study relatively large problems, compared to, for example, the quantum convolutional neural network (QCNN) [53], as quantum computation is required only for a small subroutine of the full forward pass. The models process data mostly classically and use only small circuits on local patches of a data point. Details of the model can be found in Sec. IV B.

III. DATASETS

Topologies seen in LArTPC images fall into two main categories: tracks and showers. Tracks are linear structures, whereas showers are dense collections of nonzero pixels. Showers are produced by electromagnetic cascades induced by electrons (e^-), positrons (e^+) and photons (γ), whereas other particles such as muons (μ^\pm), charged pions (π^\pm) and protons (p) produce tracks. Clear examples of tracks and showers can be seen in Fig. 3.

This study uses two LArTPC datasets to test the efficacy of QML in hit-based topology classification. The first is the openly available MicroBooNE dataset [1,54], which is described in detail in Sec. III A. The second is an original dataset produced specifically for this study. It is made up of events where a muon and an electron are produced at the same vertex but at various opening angles. It allows us to easily control the “difficulty” of classification of the events produced. Intuitively, the smaller the angle between the two particles, the “denser” the events created are. Details of this dataset can be found in Sec. III B.

A. MicroBooNE open dataset

This dataset, utilizing the Booster Neutrino Beam (BNB) flux [55], consists of approximately 750,000 events simulated in the MicroBooNE detector, featuring neutrino interactions producing electrons, photons, muons, charged pions, and protons and overlaid real cosmic rays. The MicroBooNE experiment is a LArTPC neutrino detector, which operated as part of Fermilab’s SBN program. The detector is positioned on the surface, which results in a significant background of cosmic-ray-induced particles, predominantly muons. Due to its long integration time, MicroBooNE detector collects ionization over an extended drift length, further amplifying cosmic-ray contamination. However, these cosmic-ray-induced particles have been removed for this study to ensure that the analysis focuses on

distinguishing neutrino-induced track and shower events. Moreover, this is to align the study more closely with DUNE, where the cosmic ray background will be negligible. No other cuts or modifications were applied to the dataset.

Each event in the dataset is represented as a set of three two-dimensional images, one for each wire readout plane: two induction planes (U and V), each consisting of 2400 wires, and a collection plane (Y) with 3456 wires. The induction planes are oriented at $\pm 60^\circ$ relative to the collection plane, whose wires run vertically. The wire spacing is 3 mm for all planes. In each image, the x-axis corresponds to the wire number, while the y-axis represents the drift time. Pixel intensity is proportional to the number of ionized electrons reaching that position and time, effectively encoding energy deposition information. The effective pixel resolution is 3.3 mm in the drift (y) direction and 3 mm in the wire (x) direction, as the waveform is integrated over six TPC time-ticks (3 μ s).

The goal of this work is to separate track and shower topologies. To facilitate training, we subdivide event images into $N \times N$ pixel patches (See Fig. 2) and train on randomly shuffled patches across all the events. Mimicking the approach used in another study [26], a patch around each hit is built, with the chosen hit in the center of the patch. Features of that patch are analyzed to determine whether the central pixel is track or showerlike. By repeating this process for each hit in the image, the algorithm is able to characterize each hit in the image as track or showerlike. Based on the simulation labels, each patch is assigned a label based on the particle that contributes most to the intensity of the central pixel. While multiple charged particles can deposit charge at the same location, we adopt a single-particle labeling scheme, which has been shown to be beneficial for downstream tasks [56]. At this initial stage of reconstruction, rather than identifying the specific particle responsible for the signal, each deposit is classified as tracklike (label 0) or showerlike (label 1). This labelling scheme is well-motivated, as each charged particle produces either a tracklike or showerlike signature.

B. Neutrinolike dataset

The aim of this dataset is to create events with one showerlike and one tracklike particle emanating from a common vertex with a variable opening angle between them. This was performed using LArSIMple, a simple LArTPC detector simulation [57] based on Geant4 v4_11_1_p01ba [58]. An electron was created with its direction randomly chosen from an isotropic distribution, and a muon was created within a cone of a variable opening angle around the direction of the electron. The produced muons and electrons have energies drawn from a flat distribution between 1.0 GeV and 5.0 GeV. Together, these two particles produce interactions topologies with one showerlike and one tracklike object.




FIG. 2. Randomly selected example event patches at various image sizes, categorized by pixel-level track and shower labels. Each row corresponds to a fixed patch size (given in pixels on the left), while each column shows the same event at different patch size scales. The label is determined by the classification of the central pixel as either tracklike or showerlike.

The particles are tracked through a cuboid detector filled with liquid argon and the dimensions in the (x, y, z) directions are $5\text{ m} \times 5\text{ m} \times 5\text{ m}$, where z defines the beam direction, y is vertical and x is the drift direction. The simulation produces three-dimensional energy deposits within the detector volume that are converted into three images with height coming from the drift coordinate x and width from one-dimensional projections of the yz plane, similar to the three wire readout planes in the planned DUNE detectors [59]. These three views are referred to as u , v and w and are aligned at 35.9° , -35.9° and 0° to the vertical, respectively.

Three example events are shown in Fig. 3 with opening angles of 0° , 10° and 25° between the electron and muon, as seen in one of the three readout planes in the coordinates

(wire, x). The showerlike activity in the events is shown in orange, and the tracklike deposits are shown in blue. It is clear that as the opening angle increases, the amount of overlap between the track- and showerlike energy deposits decreases.

IV. MODEL DESIGN CONCEPTS

The models used in this study are all based on convolutional neural network (CNN) architectures, which have achieved remarkable success in computer vision and beyond. CNNs form the backbone of historic architectures such as AlexNet (2012) [60], ResNet (2015) [61] and UNet (2015) [62]. They also play a key role in modern generative models, such as stable diffusion (2022) [63].




FIG. 3. Three example events containing one electron and one muon produced at a common vertex position. The events are seen projected into one of the three readout views, and the projected opening angle is given. Showerlike energy deposits in the electron-initiated electromagnetic shower and the small ionization electrons along the muon track are shown in orange, and the energy deposits coming directly from the muon are shown in blue.




FIG. 4. A convolutional neural network and its extension to a quanvolution network. Data (grayscale) can be processed classically (red) or using quantum circuits (blue) where indicated. Quantum circuits can be used as a replacement for dot products with classical filters or classical classifier networks (See Sec. IV and Appendix B). M indicates a measurement whose expectation value is the output of the circuit.

The convolutional layer uses a sliding filter which processes local windows of an image, sharing the same weights for all windows. This way it can effectively extract crucial features from the image. Given that LArTPC events can be interpreted as a collection of images, as described in Sec. III, CNNs are naturally suited for their analysis. Numerous existing studies have successfully applied CNNs to LArTPC data [1,21,22,25,28,64,65].

One of the main properties of a convolutional layer is translation equivariance; the property of “respecting” translation symmetries of the input data. Section IV A explores this property in more detail and discusses how convolutions can be extended to respect symmetries beyond translation.

In this study, quantum circuits are introduced as replacements for specific components of a standard CNN architecture, resulting in a quanvolutional network. The quanvolutional layer, detailed in Sec. IV B, modifies classical convolution while inheriting the symmetries of any convolutional layer. After a series of convolutional or quanvolutional operations, the generated feature maps are passed to a classifier, which may be either classical [a multilayer perceptron (MLP)] or quantum [a parametrized

quantum circuit (PQC)]. Section IV C discusses the use of an equivariant quantum classifier for this task. Pooling layers may be introduced when dimensionality reduction is needed (See Fig. 4).

A. Convolution and its equivariance

Equivariance is a property of a map $\mathcal{M}: X \rightarrow Y$ between two spaces (formally, G-sets) satisfied when [66]

$$\mathcal{M}(gx) = g\mathcal{M}(x) \quad \forall g \in G, \quad \forall x \in X. \quad (1)$$

This can be interpreted as: acting with a group element before acting with the map results in the same object in Y as first mapping to Y and then acting with the group element (See Fig. 5).

The convolution layers of a CNN are translation-equivariant maps. This means that if the original input becomes shifted, the output of the layer will be shifted accordingly. Mathematically, a convolution² can be expressed as [67,68]

$$[F \star W](g) = \sum_{h \in G} W(g^{-1}h)F(h), \quad (2)$$

where $g \in G_{\text{out}}$ are elements of the group the output is defined over, $h \in G$ are the elements of the group the input is defined over (can be different to G_{out}), W is the filter and F is the feature map (original data point in the case of the first convolution layer). Having defined the convolution this way, we can express the most common 2D convolution




FIG. 5. Equivariant map \mathcal{M} and an element g of a group G acting on elements x, y of two spaces X, Y , respectively.

²Technically, this is a closely related function called cross-correlation, but the term convolution is prevalent in ML literature.

by taking $G = G_{\text{out}} = (\mathbb{Z}^2, +)$, the (pixelized) translation group.

For tasks dealing with images, of interest are convolutions defined over roto-translation groups; subgroups of the Euclidean group $E(2)$ [69]. The Euclidean group can be expressed as a semidirect product of the translation group $(\mathbb{R}^2, +)$ and the orthogonal group $O(2)$ which includes all continuous translations and reflections. For many image datasets, their crucial features do not depend on the orientation of the object in a picture. Here, we hypothesize that orientation does not play a role in the topology of a track or shower.

For this study, we shrink the $O(2)$ group to just the C_4 group (the four 90° rotations) for simplicity. This results in an overall $(\mathbb{Z}^2, +) \rtimes C_4$ group of interest. This can be implemented by having each pixel contain a vector of four values, one for each element of C_4 , and correctly applying group elements to the filters in Eq. (2).

We note that this is the first investigation into the use of Euclidean symmetries on LArTPC images.

B. Quanvolution and its equivariance

The quanvolutional neural network has been proposed in [70,71] as a quantum-enhanced extension to the convolutional neural network.

We first note a common misinterpretation of the quanvolution: this architecture is often presented as using quantum circuits as “quantum filters”³ performing a convolution [33,70–72]. Appendix B shows that this interpretation is misleading and highlights a more subtle but crucial relationship between the quanvolution and convolution. This new interpretation can be used to show how quanvolutions obtain their equivariance property.

The quanvolution layer takes, as input, a feature map F . Windows of the feature map (like those in a convolution) are embedded in a unitary $U_{F,W}$ on the Hilbert space of a PQC. The number of qubits n_Q used is ordinarily equal to the number of pixels in the window. The full circuit processes the pixels contained in the window and returns a real number as output. This is usually obtained via an expectation value of a chosen observable M . This real number is used to define the feature map of the following layer. The feature extraction part of Fig. 4 represents this visually. In this work, the quanvolution circuits used are reuploader circuits of the form

$$F_{\text{out}}(g) = \langle \psi_0 | (U_{F,W}^\dagger U_\theta^\dagger(g))^R | M | (U_\theta U_{F,W}(g))^R | \psi_0 \rangle, \quad (3)$$

where $|\psi_0\rangle$ is some initial state of the qubits, taken in this study to be the $|0\rangle^{\otimes n_Q}$, U_θ is an ansatz—a unitary with trainable parameters and R is the number of times the $U_\theta U_{F,W}(g)$ -block is repeated. A general circuit of this kind is pictured in Fig. 6.




FIG. 6. General n_Q -qubit reuploader circuit architecture used for all circuits in this work. $U_F(x)$ is a data-encoding unitary for a datapoint x . U_θ is a trainable ansatz which does not depend on data. M denotes a chosen measurement unitary whose expectation value is used as the output of the trainable circuit. Exact choices for $U_F(x)$, U_θ , M are discussed in the relevant sections of the text.

$U_{F,W}$ used was the rotation embedding

$$U_{F,W}(g) = \prod_{i \in [n_Q]} RX_i(F_{W(g),i}), \quad (4)$$

where RP denotes a Pauli P rotation, $P \in X, Y, Z$ and $F_{W(g),i}$ being the value of the feature map (pixel) assigned to qubit i in the window of g . U_θ was a nearest-neighbor entangling ansatz with two trainable parameters per qubit, per entanglement block layer,

$$U_\theta = \prod_{l \in [L]} RZ_{n_Q}(\theta_{2,l,n_Q}) RY_{n_Q}(\theta_{1,l,n_Q}) \times \prod_{i \in [n_Q-1]} CX_{i,i+1} RZ_i(\theta_{2,l,i}) RY_i(\theta_{1,l,i}), \quad (5)$$

where $CX_{i,j}$ is a controlled X (also called a controlled NOT) gate with i being the control and j the controlled qubit and L is a hyperparameter denoting the number of layers of the ansatz (See Table II).

A G -equivariant quanvolution can be achieved in analogy to the standard quanvolution. That is, every dot product between a feature map and a filter of the equivariant convolution described in Sec. IV A is replaced with a pass of the feature map through a quantum circuit, as described earlier in this section. We note that the equivariance of this layer comes purely from its classical construction and the use of identical circuit structure for each group element g . Appendix B shows how a quanvolution inherits its equivariance from a convolution. In other words, the individual circuits used do not need to be equivariant maps. This is unlike what will be discussed in Sec. IV C, where the invariant quantum classifier is constructed using methods from geometric quantum machine learning.

As mentioned earlier, it is customary to use quanvolution circuits with the same number of qubits as the number of pixels in the sliding window of the quanvolution. There is, however, no need to adhere to this restriction. As long as the quanvolution circuits used depend on a given group element solely via their embedding, any size circuit can be used. Growing the circuit to many tens of qubits will likely be necessary for it to be intractable classically.

Additionally, while small filter sizes are commonly used in CNNs, some notable results from the literature show that

³A convolution filter is often also referred to as a *kernel*.

large filters can achieve competitive performance [73,74]. Future studies should investigate increasing the size of the sliding window in quanvolutions. Depending on the ansatz used, this could quickly become intractable to known classical simulation techniques.

C. Invariant quantum classifier

Equation (3) is, at its core, a description of a quantum regression model. If we forget about the group structure required for quanvolutions, we can rewrite it more generally, for any input data point x as

$$f(x) = \langle \psi_0 | (U_x^\dagger U_\theta^\dagger)^R | M | (U_\theta U_x^\dagger)^R | \psi_0 \rangle. \quad (6)$$

The output such a model is a real number, which can be turned into a binary label by choosing a threshold, e.g., 0. This is how the classification part of Fig. 4 can be realized by a quantum circuit.

Geometric quantum learning has bloomed in the recent years. Theoretical results show that some geometric models have desirable properties [75] and have been proposed as a strong candidate for potential quantum speedup [37]. Given a group, finding an equivariant ansatz can be performed in a variety of ways [48]. A quantum classifier acting on classical data can be made *invariant* if the embedding, the ansatz and the measurement are equivariant [47,48].

The classifier circuits in this work take as input the last layer of the feature extraction part of the network (See Fig. 4). Preserving rotation equivariance between the two parts of the network can be ensured by averaging along the C_4 dimension as well as along the regular convolution channels, resulting in a single image representing all the information contained in the last feature extraction layer. Other schemes for this procedure might be possible and could be explored in the future.

The invariant circuits used as classifiers in the study also follow the general structure of Fig. 6. U_x is chosen again to be the rotation embedding described in Eq. (4), this time with the whole feature map image being passed (no sliding window necessary). This embedding is equivariant to C_4 rotations. The ansatz U_θ is constructed using equivariant gates according to a scheme detailed in Appendix D and M is a Z -basis measurement of the central qubit, also equivariant to C_4 .

Given the ansatz and hyperparameters used, the number of parameters in this layer is much smaller (on the order of 10) than those used by the MLPs of the other models (on the order of 10^3).

V. MODEL ARCHITECTURES

This section details how the concepts from the previous section have been combined to design the specific architectures used on the LArTPC datasets.

TABLE I. Summary of the architectures used in the study using the nomenclature introduced in the text. Number of parameters is given as an order of magnitude estimate as the choice of hyperparameters impacts the exact number. Top three models have no symmetries beyond translations. The remaining are equivariant to C_4 rotations at least in the feature extraction part.

Model	Feature extraction	Classification	No. params
<i>Deep CNN</i>	<i>CNE</i>	<i>CNE</i>	10^5
<i>CNN</i>	<i>CNE</i>	<i>CNE</i>	10^3
<i>Quanv</i>	<i>QNE</i>	<i>CNE</i>	10^3
<i>Deep GCNN</i>	<i>CE</i>	<i>CNE</i>	10^5
<i>GQuanv</i>	<i>QE</i>	<i>CNE</i>	10^3
<i>InvQuanv</i>	<i>QE</i>	<i>QE</i>	10^2

We consider the CNN to be made up of two modules: the feature extraction module (with optional pooling after each layer) and the classification module. To help keep track of the different variants of the models, we refer to them by the symmetry (equivariance) status S and type T of their modules, where $S \in \{E, NE\}$ and $T \in \{Q, C\}$. E denotes equivariant, NE —nonequivariant, Q —quantum, C —classical. Together, a given architecture can be referred to as $TS \rightarrow TS$. Table I details all the models used.

A. Classical benchmarks

Three classical benchmark models are used in this study:

- (i) *Deep CNN* ($CNE \rightarrow CNE$)—A six-layer convolutional neural network with 3×3 filters and output channels [16, 16, 32, 32, 64, 64], in each layer respectively.
- (ii) *CNN* ($CNE \rightarrow CNE$)—A compact convolutional network with 2 layers of 2×2 filters and three output channels per layer.
- (iii) *Deep GCNN* ($CE \rightarrow CNE$)—A deep CNN that incorporates C_4 group convolutions, introducing rotational symmetry while maintaining the same overall architecture as the deep CNN.

All models use an MLP with two hidden layers as the final classifier. The shape of the hidden layers was a hyperparameter (See Appendix C)

B. Quantum-enhanced architectures

Three quanvolutional networks with 2×2 quanvolution windows are implemented:

- (i) *Quanv* $QNE \rightarrow CNE$ —Two quanvolutional layers, each with three output channels, followed by an MLP.
- (ii) *GQuanv* $QE \rightarrow CNE$ —Two C_4 -symmetric quanvolutional layers, each with three output channels. Classification with an MLP.
- (iii) *InvQuanv* $QE \rightarrow QE$ —Two C_4 -symmetric quanvolutional layers, each with three output channels.

Classification with a C_4 -symmetric quantum classifier using an equivariant embedding, equivariant reupload circuits and an equivariant measurement.

Quanv is a standard quanvolutional network, utilizing only the $(Z_2, +)$ group. The other two are rotation-aware. Feature extraction in both is achieved with a C_4 -symmetric quanvolution described in Sec. IV B. They differ only in their classifier module. The first uses a classical MLP and the other, a quantum, symmetry-aware classifier discussed in Sec. IV C. The second model is a fully *invariant* classifier. This means that two images which differ by a C_4 rotation will result in the same output, leading necessarily to the same label.

VI. RESULTS

For clarity, all plots in this section contain only the best performing quantum, classical and deep classical models. Figures 10–12 in Appendix A show the performance of all models discussed.

A. Performance on the MicroBooNE open dataset

In this section, we evaluate the performance of our proposed models on the open MicroBooNE dataset. Specifically, we investigate how the spatial context surrounding a classified pixel influences the model’s performance. Figure 2 illustrates examples of the same pixel within an event as part of progressively larger patches.

In our framework, a patch refers to the local neighborhood of a pixel that is provided as input to the model. While the model ultimately classifies only the central pixel (as either “track” or “shower”), it utilizes the surrounding information within the patch to make this determination. For instance, if a pixel lies within a straight, connected path and exhibits consistent energy signatures, the model is likely to classify it as part of a track.

Incorporating a larger field of view offers contextual information about the broader topology, aiding classification. However, excessively large patches may introduce unrelated or irrelevant structures, potentially degrading model performance by introducing noise. The optimal patch size for LArTPC pixel classification remains an open question. In this work, we contribute to this ongoing investigation by evaluating model performance on the simplified MicroBooNE dataset described in Sec. III A.

Figure 7 presents the performance of the best performing quantum, classical and deep classical architectures described in Sec. IV across different patch sizes. Notably, larger models exhibit the ability to process increasingly large patches effectively—the more contextual information they receive, the better their classification performance. Conversely, smaller models demonstrate a decline in performance when presented with patches that are too large, likely due to an inability to process the additional information effectively.

Quantum models demonstrate superior performance compared to their classical counterparts across multiple




FIG. 7. Receiver operating characteristic area under curve (ROC AUC) for best performing architectures used in this study with increasing pixel patch size. The deep model was not defined for small patches as no padding was used and it would reduce the size of feature maps to 1×1 before classification. Models were trained on 500 patches.

patch sizes when constrained to a similar number of trainable parameters and with an identical number of filters. A 100-fold increase in the number of parameters allows classical models to surpass the performance of the quantum-enhanced architectures.

It is unclear whether the inclusion of C_4 geometry in the models can be of undeniable benefit. Although among the deep models, the symmetric model showcased a consistently slightly better performance, among the quantum models, this bias shows no real advantage. A study incorporating larger rotation groups into the models is paramount.

A key limitation of large-scale studies involving the simulation of quantum devices on classical hardware is the substantial computational cost in terms of time and RAM. In our study, performing hyperparameter optimization for selecting the best model (details in Appendix C) at each patch size (Fig. 7) requires multiple days per patch size. This computational overhead imposes constraints on the volume of training data and number of filters that can be incorporated into the optimization process.

To further investigate the learning capabilities of our models, we analyze their performance as a function of training dataset size. Specifically, we evaluate the best-performing hyperparameter configurations for models found in the previous analysis. Using a patch size of 21, we check their performance across training set sizes ranging from 100 to 1000 data points. The metrics for each model are reported by averaging the 10 best-performing instances from a run, using hundreds of random seeds for each training size. We report the mean, standard deviation, and maximum performance. The primary goal of this study is to assess the extent to which models can generalize in a data-constrained setting.

The results of this experiment are presented in Fig. 8. The overall trend confirms the hierarchy of the quantum,




FIG. 8. Learning capability of models used in this study. Vertical axis is the receiver operating characteristic area under curve (ROC AUC) and the horizontal axis is a training size (number of patches). Patch size was kept at 21.

classical and deep classical architectures in the low-data regime we are able to probe.

The greater performance of the quanvolutional NN compared to the shallow CNN suggests that the representations learned by the quanvolutional NN are richer and more informative. This enables better feature extraction for the final classification layers, a trend which remains even as we increase the amount of training data. This suggests that (classical) quantum-inspired or actual quantum-enhanced architectures may offer competitive feature representations. However, further investigations into their scalability with larger datasets are warranted.

B. Performance on the neutrino-like dataset

As described in Sec. III B, events in this dataset are characterized by a simple yet informative parameter—the angle between the two created particles. This angle serves as a proxy for the classification difficulty of individual pixels: smaller angles correspond to more challenging classification tasks due to increased spatial overlap and denser regions of energy deposition.

Understanding how models perform in high-density regions is crucial for future LArTPC experiments, as these regions present a major bottleneck for accurate event reconstruction. To systematically investigate how model performance scales with complexity, we divide the dataset into three categories based on the inter-particle angle:

- (i) *Easy*: 15° – 180° (widely separated particles)
- (ii) *Medium*: 5° – 15° (moderate overlap)
- (iii) *Hard*: 0° – 5° (significant overlap and the most complex topology)

We note that these bins were chosen empirically based on visual inspection of the events not on a detailed examination of the complexity of individual patches. One could devise more precise metrics by quantifying the amount of “shower” pixels in a “track” patch, for example. Results in Sec. VI B confirm the average-case




FIG. 9. Performance in different environments of selected models in this study. Vertical axis is the receiver operating characteristic area under curve (ROC AUC) and the horizontal axis is the opening angle between the electron and muon particles. The patch size was kept at 21.

hardness of the three datasets created. The performance of the models across these three categories is summarized in Fig. 9. A clear pattern emerges when comparing the shallow CNN to the quanvolutional NN: while the quanvolutional NN consistently outperforms the shallow CNN across all levels of difficulty, the largest performance gap appears in the hardest category (0° – 5°). This suggests that the quanvolutional NN is better equipped to resolve complex topologies and extract meaningful features, a crucial property when dealing with busy detector regions. We can potentially attribute this to the richer feature maps produced by the quanvolutional NN.

However, we also observe that larger classical CNNs—with significantly more filters and overall capacity—continue to outperform the quanvolutional NN across all difficulty levels. Interestingly, it is the nonequivariant architecture which, in general, prevails. While each of the events in this dataset has a sense of direction (most patches will have structures oriented along the direction of the muon), which could render the addition of symmetries unhelpful, the patches used for training and testing are sampled from 100 different events, where that direction is isotropic.

VII. CONCLUSION

This paper contributes to the fields of experimental neutrino physics and quantum machine learning. In the context of particle physics, we address crucial open questions in ML-driven pixel-level LArTPC classification. We examine the effects of patch size on the performance of the models, the benefits of including rotational symmetries in their design and begin to study their behavior in dense environments. In the quantum machine learning domain, we point out a common misconception about a popular model, define it in a rigorous way based on its connection to the convolution layer, and show, for the first time, how to design it such that it respects symmetries beyond translation.

Results show that quantum-enhanced models are competitive with fully classical ones. For a similar number of parameters, quantum models consistently achieve better results. They remain outperformed by classical models with 10 times more parameters, however. This trend holds for all the analyses implemented within this work.

Using the MicroBooNE data, it has been shown that large models can utilize large patches for pixel classification, while the amount of information contained in these large patches becomes confusing for smaller models. These find optimal performance at a patch size of 21. An obvious question follows: can the large models benefit from even larger patches?

One downside of our study is that we operate in a very low-data regime, showing results for models trained on up to only 1000 samples. Given the amount of data available in LArTPC experiments and HEP experiments in general, the use of more efficient QML pipeline techniques becomes paramount. Future studies should strive to reduce the computational training and test time by leveraging GPUs, just-in-time compilation or simulations beyond state-vector-methods.

Implementing a custom “neutrinolike” dataset, we defined a single parameter (angle between the two produced particles) as a proxy to denote the difficulty of classifying patches from a given event. Recognizing a similar way to categorize pixels in actual experiments could become a highly beneficial tool for analyzing the performance of the implemented classifiers.

The inclusion of C_4 symmetry in the models does not guarantee a substantial advantage. No advantage is seen within the (small) quanvolution models—the nonsymmetric quanvolution performs best in almost all scenarios. A small advantage of the (large) symmetric classical model over its nonsymmetric counterpart can be seen in the MicroBooNE dataset but not in the neutrinolike dataset. A study varying the symmetry group (adding reflections, probing C_n with $n \neq 4$, or going to the continuous $SO(2)$ regime) is needed to determine which (if any) provides the best inductive bias for models dealing with LArTPC images. The case for discrete groups in quanvolutions can be explored using methods described in this work.

An obvious direction for future study is an extension to larger subgroups of the Euclidean group. We note also that some potentially interesting combinations of the feature extraction and classification modules have not been explored. Namely, $CE \rightarrow CE$, $QE \rightarrow CE$ and $CE \rightarrow QE$. These would be fully invariant classifiers. Contrasting the properties and performance of all the combinations of the feature extraction and classification modules could further our understanding of any potential advantages offered by quantum elements in the architecture, as well as the inclusion of symmetry in the model.

Another interesting continuation of this study is the separation of similar showerlike deposits coming from different particles. In dense regions, different interactions can produce showerlike topologies that are very close to each other or even overlap, as is the case with the decay

photons from π^0 mesons. As it remains crucial for a correct reconstruction to separate each deposit at the pixel level, work in that direction would likely be beneficial.

A QML approach, such as the one proposed for this work, could be in the future implemented inside pattern recognition tools such as Pandora. Given Pandora’s modular and multialgorithmic nature, QML algorithms could be used alongside classical CNNs. Depending on the performance, QML could be deployed to solve a particular problem where classical computing cannot. This would be particularly valuable with data coming from big detectors such as the DUNE-FD, where we expect a very high level of complexity and a hybrid approach in reconstruction could be very valuable.

ACKNOWLEDGMENTS

M. J. thanks Peter T J Bradshaw for valuable discussions. We acknowledge the MicroBooNE Collaboration for making publicly available the datasets [54] employed in this work. These data sets consist of simulated neutrino interactions from the Booster Neutrino Beamline overlaid on top of cosmic data collected with the MicroBooNE detector [1]. The datasets were produced using the MicroBooNE Genie Tune [76]. C. D. and M. J. contributed equally to this work. The Pandora software project is supported by the Science and Technology Facilities Council (STFC), part of the United Kingdom Research and Innovation.

DATA AVAILABILITY

The data that support the findings of this Letter are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: RESULTS FOR ALL MODELS

This section contains Figs. 10–12, which are the full versions of Figs. 7–9.




FIG. 10. Receiver operating characteristic area under curve (ROC AUC) for all models used in this study with increasing pixel patch size. The deep models were not defined for small patches as no padding was used and it would reduce the size of feature maps to 1×1 before classification.




FIG. 11. Learning capability of all models used in this study. Vertical axis is the receiver operating characteristic area under curve (ROC AUC) and the horizontal axis is a training size (number of patches). Patch size was kept at 21.




FIG. 12. Performance in different environments of all models used in this study. Vertical axis is the receiver operating characteristic area under curve (ROC AUC) and the horizontal axis is the opening angle between the electron and muon particles. The patch size was kept at 21.

APPENDIX B: THE MISINTERPRETED QUANVOLUTION

As hinted at in Sec. IV, the quanvolution layer used in this study does not—contrary to common interpretation—implement a quantum circuit as a “filter” in an otherwise classical convolution.

Recalling Eq. (2), we note that the convolution involves a dot product of two functions defined on a common space of group elements, where one of the functions is “moved” through that space.

In order for the quantum circuit to act as the filter, it would have to be defined on the pixels of the feature map [68]. One should be able to ask “What is the value of the quantum circuit at pixel (2,2)?”. From the construction of the quanvolution defined in the paper it is clear that such question is meaningless without first finding the value of the feature map at that point—the “value of the circuit” is

necessarily defined through the feature map which becomes embedded into it.

One *can* think of the quantum circuit as a filter in the image-processing sense, a predefined operation applied to a neighborhood of a pixel [77]. This should not, however, be extended to thinking of the circuit as the filter of a convolution, based on the above argument.

1. Where is the convolution?

As reasoned above, the circuit itself cannot be thought of as a filter in a convolution. The quanvolution, however *feels* like a convolution; it seems to inherit its equivariance property, has a stride, a window size and can be implemented, algorithmically, in a similar way to a convolution. It also, as shown in the main text, seems to very naturally extend to group-quanvolutions based on group-convolutions. One way of understanding the connection between the two is to consider how the feature map of a layer is embedded into the quantum circuit.

Let us consider a convolution filter W which is *matrix-valued*. Observe now that if these matrices $W(g)$ are restricted to being Hermitian and commuting with each other, the embedding can be expressed as

$$U_{F,W}(g) = \exp \left(-i \sum_h F(h) W(g^{-1}h) \right) \\ = \exp (-i[F \star W](g)). \quad (\text{B1})$$

The embedding circuit is *generated* by a convolution between the feature map F and a “quantum filter” W . Figure 13 demonstrates this idea. The simple rotation embedding used in the main text [Eq. (4)] can be achieved with a filter containing one of the $\{X_i | i \in [n_Q]\}$ gates at each entry.




FIG. 13. The data-embedding unitary pictured as a convolution with a Hermitian-matrix-valued filter.




FIG. 14. A C_4 -invariant gate obtained from twirling X_3Y_6 overlaid on a pixelized grid representing the input layer to the classifier defined in the text.

With Eq. (B1), the quanvolution operation can be written as

$$\begin{aligned} [F \star^Q W](g) &= \text{Circ}(U_{F,W}(g)) \\ &= \text{Circ}(\exp(-i[F \star W](g))), \end{aligned} \quad (\text{B2})$$

where the $\text{Circ}(U(g))$ function obtains a real-valued property of a state which depends on g only through $U_{F,W}(g)$. An example would be a reuploader circuit utilizing expectation values, like the one used in the main text

$$\begin{aligned} \text{Circ}(U_{F,W}(g)) &= \langle \psi_0 | (U_{\theta}^{\dagger} U_{F,W}^{\dagger}(g))^R | M | (U_{\theta} U_{F,W}(g))^R | \psi_0 \rangle. \end{aligned} \quad (\text{B3})$$

An example of a circuit which does not satisfy the above constraint would be

$$\begin{aligned} \text{Circ}(U_{F,W}(g)) &= \langle \psi_0 | (U_{\theta}^{\dagger} U_{F,W}^{\dagger}(g))^R | M_g | (U_{\theta} U_{F,W}(g))^R | \psi_0 \rangle, \end{aligned} \quad (\text{B4})$$

where the observable M depends on the position in the produced feature map.

To illustrate how the quanvolution inherits the equivariance property from a convolution without being a convolution itself, let us begin with an important property of feature maps in G -convolutions

$$gF(x) = F(g^{-1}x). \quad (\text{B5})$$

This simply states that to access a value of the g -transformed feature map one can look up that value in the original feature map at the point which transforms to x via g . Now, we show a well-known proof for the equivariance of a G -convolution [68,78]

$$\begin{aligned} [uF \star W](g) &= \sum_{h \in G} F(u^{-1}h) W(g^{-1}h) \\ &= \sum_{h \in G} F(h) W(g^{-1}uh), (h \rightarrow uh) \\ &= \sum_{h \in G} F(h) W((u^{-1}g)^{-1}h) \\ &= [F \star W](u^{-1}g) \\ &= u[F \star W](g), \end{aligned} \quad (\text{B6})$$

where the last line holds as the convolution defines a new feature map.

Now for a quanvolution

$$\begin{aligned} [uF \star^Q W](g) &= \text{Circ}(U_{uF,W}(g)) \\ &= \text{Circ}(\exp(-i[uF \star W](g))) \\ &= \text{Circ}(\exp(-i[F \star W](u^{-1}g))) \\ &= [F \star^Q W](u^{-1}g) \\ &= u[F \star^Q W](g), \end{aligned} \quad (\text{B7})$$

where the third line comes from (B6), and again in the last line, we used the property of feature maps (B5). This illustrates how any G -quanvolution can be obtained from an appropriate G -convolution and explains why the proposed C_4 -quanvolution works.

APPENDIX C: DETAILS OF HYPERPARAMETER OPTIMIZATION

For training and hyperparameter optimization, RAY [79] and WEIGHTSANDBIASES (WANDB) [80] packages were used extensively.

In the patch size study, for each of the patch sizes, a hyperparameter search over around 200 models per architecture has been performed. Models were allowed to train for up to 100 epochs but were often cut short by the RayScheduler object used in the pipeline. The RayScheduler class contains methods for speeding up wide model searches by cutting jobs short based on a metric of choice. For this study, models were terminated based on the loss they achieved during training.

Loss, validation loss and validation accuracy were reported throughout training to online and local WANDB project directories, which helped with visual assessing of their performance and accessing trained models. After training, the model used for testing was chosen based on the highest validation accuracy achieved during training.

APPENDIX D: DETAILS OF THE QUANTUM GEOMETRIC CLASSIFIER

The geometric classifier circuits are invariant to rotations of the feature maps received by them (coming from the last layer of the quanvolution and potential pooling). These classifier circuits, for ease of software implementation, were defined only on 9-qubit circuits. That means that only 3×3 patches could be classified. This provided a restriction for the feature extraction layers (See Fig. 4) preceding the quantum classifier. That is, the combination of the quanvolutions and the classical pooling had to result in a final layer that has a shape 3×3 (before being flattened).

TABLE II. Hyperparameters optimized for the models used in the study. $[]$ signifies a continuous range and $\{\}$ a discrete set.

Hyperparameter	Sampled set	Models
Learning rate	$[10^{-3}, 10^{-1}]$	All
Number of layers (quanvolution circuits)	$\{1, 2\}$	QECNE, QEQUE, QNECNE
Number of filters (quanvolution)	$\{1, 2, 3\}$	QECNE, QEQUE, QNECNE
Number of reuploads (quanvolution circuits)	$\{1, 2\}$	QECNE, QEQUE, QNECNE
Maximum for parameter initialization (quanvolution circuits)	$\{10^{-3}, 10^{-1}, \frac{\pi}{4}, \frac{\pi}{2}, 2\pi\}$	QECNE, QEQUE, QNECNE
Dense units sizes	$\{(8, 8), (128, 32)\}$	QECNE, QNECNE
Use of dropout	$\{True, False\}$	QECNE, QNECNE
Dropout amount	$\{10^{-4}, 0.5\}$	QECNE, QNECNE
Classifier number of layers	$\{1, 2, 3, 4\}$	QEQUE
Classifier number of reuploads	$\{1, 2, 3\}$	QEQUE
Number of 1-local gates	$\{1, 2, 3, 4\}$	QEQUE
Number of 2-local gates	$\{1, 2, 3, 4\}$	QEQUE
Placement of 1-local gates	See Appendix D	QEQUE
Placement of 2-local gates	See Appendix D	QEQUE

The invariant classifiers' layers have been obtained according to the following protocol:

- (i) randomly pick a number N_1 of single-qubit gates (See Table II) (e.g., 1)
- (ii) randomly pick a number N_2 of two-qubit gates (See Table II) (e.g., 1)
- (iii) randomly pick a qubit for each of the 1-qubit gates to act on (e.g., qubit 3)
- (iv) randomly pick a pair of qubits for each of the 2-qubit gates to act on (e.g., qubits 3 and 6).
- (v) for each of the gates, randomly choose a Pauli string of the correct length (e.g., X , XY), creating generators (X_3 , X_3Y_6)
- (vi) by twirling the generators (see an example in Fig. 14), obtain generators of C_4 -equivariant gates ($X_1 + X_3 + X_5 + X_7$, $X_3Y_6 + X_1Y_0 + X_5Y_2 + X_7Y_8$)

- (vii) reject the choice if the resulting generators cannot be implemented as combinations of 2-local gates, accept otherwise

The rejection criterion in the last step refers to a situation like:

$$\tau(X_0Z_2) = X_0Z_2 + X_2Z_8 + X_8Z_6 + X_6Z_0. \quad (D1)$$

Because the elements of this generator do not commute, the resulting parametrized gate cannot be implemented as a product of 2-qubit gates, $e^{i\theta\tau(X_0Z_2)} \neq e^{i\theta X_0Z_2} e^{i\theta X_2Z_8} e^{i\theta X_8Z_6} e^{i\theta X_6Z_0}$.

We note a detailed study of rotation-invariant circuits for images in [81], where the feature extraction is achieved using equivariant convolutions (our work uses equivariant quanvolutions).

-
- [1] R. Acciarri *et al.* (MicroBooNE Collaboration), *J. Instrum.* **12**, P03011 (2017).
 - [2] C. A. Argüelles, I. Esteban, M. Hostert, K. J. Kelly, J. Kopp, P. A. N. Machado, I. Martinez-Soler, and Y. F. Perez-Gonzalez, *Phys. Rev. Lett.* **128**, 241802 (2022).
 - [3] P. A. Machado, O. Palamara, and D. W. Schmitz, *Annu. Rev. Nucl. Part. Sci.* **69**, 363 (2019).
 - [4] B. Abi *et al.* (DUNE Collaboration), *Eur. Phys. J. C* **80**, 978 (2020).
 - [5] S. Fukasawa, M. Ghosh, and O. Yasuda, *Nucl. Phys.* **B918**, 337 (2017).
 - [6] B. Brahma and A. Giri, [arXiv:2306.05258](https://arxiv.org/abs/2306.05258).
 - [7] K. Møller, A. M. Suliga, I. Tamborra, and P. B. Denton, *J. Cosmol. Astropart. Phys.* **05** (2018) 066.
 - [8] F. Domingo, H. K. Dreiner, D. Köhler, S. Nangia, and A. Shah, *J. High Energy Phys.* **05** (2024) 258.
 - [9] B. Baller, *J. Instrum.* **12**, P07010 (2017).
 - [10] P. Abratenko *et al.* (ICARUS Collaboration), *Eur. Phys. J. C* **83**, 467 (2023).
 - [11] R. Acciarri *et al.* (MicroBooNE Collaboration), *J. Instrum.* **12**, P02017 (2017).
 - [12] B. Abi *et al.* (DUNE Collaboration), *J. Instrum.* **15**, P12004 (2020).
 - [13] R. Acciarri *et al.* (SBND Collaboration), *J. Instrum.* **15**, P06033 (2020).

- [14] M. Soderberg (ArgoNeuT Collaboration), in *Meeting of the Division of Particles and Fields of the American Physical Society (DPF 2009)* (2009) [arXiv:0910.3433](#).
- [15] G. Cerati, *EPJ Web Conf.* **295**, 08012 (2024).
- [16] Fermilab, Event displays, https://microboone-exp.fnal.gov/public/approved_plots/Event_Displays.html (2021), (accessed 2025-07-15).
- [17] J. S. Marshall and M. A. Thomson, *Eur. Phys. J. C* **75**, 439 (2015).
- [18] R. Acciarri *et al.* (MicroBooNE Collaboration), *Eur. Phys. J. C* **78**, 82 (2018).
- [19] A. Abed Abud *et al.* (DUNE Collaboration), *Eur. Phys. J. C* **83**, 618 (2023).
- [20] E. L. Snider and G. Petrillo, *J. Phys. Conf. Ser.* **898**, 042057 (2017).
- [21] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **105**, 112003 (2022).
- [22] B. Abi *et al.* (DUNE Collaboration), *Phys. Rev. D* **102**, 092003 (2020).
- [23] P. Machado, H. Schulz, and J. Turner, *Phys. Rev. D* **102**, 053010 (2020).
- [24] J. Kopp, P. Machado, M. MacMahon, and I. Martinez-Soler, [arXiv:2405.15867](#).
- [25] R. Moretti, M. Rossi, M. Biassoni, A. Giachero, M. Grossi, D. Guffanti, D. Labranca, F. Terranova, and S. Vallecorsa, *Eur. Phys. J. Plus* **139**, 723 (2024).
- [26] S. Vergani (DUNE Collaboration), *Proc. Sci. NuFact2021* (2022) 179.
- [27] S. Vergani, Using cutting edge software and techniques to model and measure experimental neutrino data, Ph.D. thesis, Apollo—University of Cambridge Repository, 2024.
- [28] A. Abed Abud *et al.* (DUNE Collaboration), *Eur. Phys. J. C* **82**, 903 (2022).
- [29] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **103**, 052012 (2021).
- [30] P. Bermejo, P. Braccia, M. S. Rudolph, Z. Holmes, L. Cincio, and M. Cerezo, [arXiv:2408.12739](#).
- [31] J. Bowles, S. Ahmed, and M. Schuld, [arXiv:2403.07059](#).
- [32] S. Y.-C. Chen, T.-C. Wei, C. Zhang, H. Yu, and S. Yoo, [arXiv:2101.06189](#).
- [33] S. Y.-C. Chen, T.-C. Wei, C. Zhang, H. Yu, and S. Yoo, *Phys. Rev. Res.* **4**, 013231 (2022).
- [34] A. Delgado, D. Venegas-Vargas, A. Huynh, and K. Carroll, [arXiv:2410.12650](#).
- [35] A. Di Meglio, M. Doser, B. Frisch, D. M. Grabowska, M. Pierini, and S. Vallecorsa (CERNQTI Collaboration), *CERN Quantum Technology Initiative Strategy and Roadmap* (2021), [10.5281/zenodo.5571809](#).
- [36] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, *Rev. Mod. Phys.* **94**, 015004 (2022).
- [37] H. Zheng, Z. Li, J. Liu, S. Strelchuk, and R. Kondor, *PRX Quantum* **4**, 020327 (2023).
- [38] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, *Phys. Rev. Res.* **2**, 033125 (2020).
- [39] B. Coyle, D. Mills, V. Danos, and E. Kashefi, *npj Quantum Inf.* **6**, 60 (2020).
- [40] J. Jäger and R. V. Krems, *Nat. Commun.* **14**, 576 (2023).
- [41] R. Sweke, E. Recio, S. Jerbi, E. Gil-Fuster, B. Fuller, J. Eisert, and J. J. Meyer, *Quantum* **9**, 1640 (2025).
- [42] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, *Nat. Commun.* **12**, 2631 (2021).
- [43] M. Cerezo, M. Larocca, D. García-Martín, N. L. Diaz, P. Braccia, E. Fontana, M. S. Rudolph, P. Bermejo, A. Ijaz, S. Thanasilp *et al.*, *Nat. Commun.* **16**, 7907 (2025).
- [44] C. Duffy, M. Hassanshah, M. Jastrzebski, and S. Malik, *Quantum Mach. Intell.* **7**, 41 (2025).
- [45] P. Duckett, G. Facini, M. Jastrzebski, S. Malik, T. Scanlon, and S. Rettie, *Phys. Rev. D* **109**, 052002 (2024).
- [46] C. Tüysüz, M. Demidik, L. Coopmans, E. Rinaldi, V. Croft, Y. Haddad, M. Rosenkranz, and K. Jansen, [arXiv:2410.16363](#).
- [47] M. Ragone, P. Braccia, Q. T. Nguyen, L. Schatzki, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, [arXiv:2210.07980](#).
- [48] Q. T. Nguyen, L. Schatzki, P. Braccia, M. Ragone, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, *PRX Quantum* **5**, 020328 (2024).
- [49] J. J. Meyer, M. Mularski, E. Gil-Fuster, A. A. Mele, F. Arzani, A. Wilms, and J. Eisert, *PRX Quantum* **4**, 010328 (2023).
- [50] A. E. Paine, V. E. Elfving, and O. Kyriienko, [arXiv:2308.01827](#).
- [51] O. Kyriienko, A. E. Paine, and V. E. Elfving, *Phys. Rev. Res.* **6**, 033291 (2024).
- [52] C. A. Williams, S. Scali, A. A. Gentile, D. Berger, and O. Kyriienko, [arXiv:2411.14259](#).
- [53] I. Cong, S. Choi, and M. D. Lukin, *Nat. Phys.* **15**, 1273 (2019).
- [54] P. Abratenko *et al.* (MicroBooNE Collaboration), [10.5281/zenodo.8370883](#) (2023).
- [55] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Phys. Rev. D* **79**, 072002 (2009).
- [56] P. Abratenko *et al.*, *J. Instrum.* **16**, P02017 (2021).
- [57] A. Chappell and L. H. Whitehead, *Eur. Phys. J. C* **82**, 1099 (2022).
- [58] S. Agostinelli *et al.* (GEANT4 Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **506**, 250 (2003).
- [59] B. Abi *et al.* (DUNE Collaboration), *J. Instrum.* **15**, T08010 (2020).
- [60] A. Krizhevsky, I. Sutskever, and G. E. Hinton, in *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*, edited by F. Pereira, C. Burges, L. Bottou, and K. Weinberger (Curran Associates, Inc., Red Hook, NY, 2012), pp. 1097–1105, [10.5555/2999134.2999257](#).
- [61] K. He, X. Zhang, S. Ren, and J. Sun, [arXiv:1512.03385](#).
- [62] O. Ronneberger, P. Fischer, and T. Brox, in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*, edited by N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi (Springer International Publishing, Cham, 2015), pp. 234–241.
- [63] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021), p. 10674.
- [64] L. Domin and K. Terao, *Phys. Rev. D* **102**, 012005 (2020).

- [65] J. Liu, J. Ott, J. Collado, B. Jargowsky, W. Wu, J. Bian, and P. Baldi (DUNE Collaboration), [arXiv:2012.06181](#).
- [66] M. Adhikari and A. Adhikari, *Basic Modern Algebra with Applications* (Springer India, New Delhi, 2014).
- [67] D. Worrall and G. Brostow, in *Proceedings of the European Conference on Computer Vision (ECCV)* (2018), pp. 567–584.
- [68] T. S. Cohen, M. Geiger, and M. Weiler, *Advances in neural information processing systems* (2019), Vol. 32.
- [69] M. Weiler and G. Cesa, *Advances in neural information processing systems* (2019), Vol. 32.
- [70] M. P. Henderson, S. Shakya, S. Pradhan, and T. Cook, *Quantum Mach. Intell.* **2** (2019).
- [71] J. Liu, K. H. Lim, K. L. Wood, W. Huang, C. Guo, and H.-L. Huang, *Sci. China Phys. Mech. Astron.* **64**, 290311 (2021).
- [72] Quvolutional neural networks, https://pennylane.ai/qml/demos/tutorial_quanvolution#quanvolutional-neural-networks, (accessed: 07-02-2025).
- [73] X. Ding, X. Zhang, J. Han, and G. Ding, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 11963–11975.
- [74] S. Liu, T. Chen, X. Chen, X. Chen, Q. Xiao, B. Wu, T. Kärkkäinen, M. Pechenizkiy, D. Mocanu, and Z. Wang, [arXiv:2207.03620](#).
- [75] L. Schatzki, M. Larocca, Q. T. Nguyen, F. Sauvage, and M. Cerezo, *npj Quantum Inf.* **10**, 12 (2024).
- [76] P. Abratenko *et al.* (MicroBooNE Collaboration), *Phys. Rev. D* **105**, 072001 (2022).
- [77] R. C. Gonzalez and R. E. Woods, *Digital Image Processing (3rd Edition)* (Prentice-Hall, Inc., USA, 2006).
- [78] T. Cohen and M. Welling, in *International Conference on Machine Learning* (PMLR, New York City, NY, 2016), pp. 2990–2999.
- [79] P. Moritz, R. Nishihara, S. Wang, A. Tumanov, R. Liaw, E. Liang, M. Elibol, Z. Yang, W. Paul, M. I. Jordan *et al.*, in *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)* (2018), pp. 561–577.
- [80] L. Biewald, Experiment tracking with weights and biases (2020), software available from wandb.com.
- [81] P. S. S. Sein, M. Cañizo, and R. Orús, *Phys. Rev. Res.* **7**, 013082 (2025).