

NuGraph2: A Graph Neural Network for Neutrino Physics Event Reconstruction

A. Aurisano and V Hewes
University of Cincinnati, Cincinnati, OH 45221, USA

G. Cerati and J. Kowalkowski
Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

C. S. Lee and W. Liao
Northwestern University, Evanston, IL 60208, USA

D. Grzenda, K. Gumpula, and X. Zhang^a
Data Science Institute, University of Chicago, Chicago, IL 60637, USA
(Dated: March 19, 2024)

Liquid Argon Time Projection Chamber (LArTPC) detector technology offers a wealth of high-resolution information on particle interactions, and leveraging that information to its full potential requires sophisticated automated reconstruction techniques. This article describes **NuGraph2**, a Graph Neural Network (GNN) for low-level reconstruction of simulated neutrino interactions in a LArTPC detector. Simulated neutrino interactions in the MicroBooNE detector geometry are described as heterogeneous graphs, with energy depositions on each detector plane forming nodes on planar subgraphs. The network utilizes a multi-head attention message-passing mechanism to perform background filtering and semantic labelling on these graph nodes, identifying those associated with the primary physics interaction with 98.0% efficiency and labelling them according to particle type with 94.9% efficiency. The network operates directly on detector observables across multiple 2D representations, but utilizes a 3D-context-aware mechanism to encourage consistency between these representations. Model inference takes 0.12 s/event on a CPU, and 0.005s/event batched on a GPU. This architecture is designed to be a general-purpose solution for particle reconstruction in neutrino physics, with the potential for deployment across a broad range of detector technologies, and offers a core convolution engine that can be leveraged for a variety of tasks beyond the two described in this article.

I. INTRODUCTION

In recent years, machine learning (ML) applications have seen increasingly widespread use in neutrino physics, and more broadly in High Energy Physics (HEP). The NOvA experiment has utilized Convolutional Neural Networks (CNNs) for neutrino flavour identification [1], prong identification [2], and cosmic filtering. CNN classification is currently used for the benchmark physics sensitivity at the next-generation Deep Underground Neutrino Experiment (DUNE) [3]; it has also been leveraged for lower-level reconstruction tasks [4]. MicroBooNE has been using CNNs for various tasks, including hit classification [5] and particle identification [6].

However, dense CNN applications such as these have several limitations – they typically operate on fixed-size inputs, which is often computationally inefficient for small events, and leads to the truncation of large events. Also, given the natural sparseness of most particle interactions, the dense representation leads to the majority of computations being wasted on empty space.

These limitations have been resolved through the adoption of sparse CNNs, utilizing tools such as MinkowskiEngine [7]. Sparse CNNs operate only on activated pixels within an infinite manifold, which resolves issues both with wasted computations and truncations. Applications in Liquid Argon Time Projection Chambers (LArTPCs) have shown great promise at particle hit identification and clustering [8–10].

Although sparse CNN applications resolve the major issues with dense CNNs, many more limitations inherent to the CNN paradigm persist. For instance, the kernel convolutions that define the approach require inputs to be fixed to a regular Euclidean grid; in the context of particle physics detectors, this can require arbitrary transformations to be applied to detector data. For instance, a 3D point cloud produced from detector hits must be aggregated into voxels in order to construct CNN inputs. Since the network output inhabits this voxelized domain, the end user must either perform this awkward transformation in reverse, or arbitrarily decide to work within the voxelized domain dictated by the choice of reconstruction algorithm.

^a also at *University of California, Los Angeles, Los Angeles, CA 90095, USA*

This domain issue can be addressed with Graph Neural Networks (GNNs). Unlike CNNs, GNNs operate on a collection of graph nodes, along with edges that define their connectivity. This much more flexible data structure allows for network architectures that can operate on the data in its native domain without any transformations required. This flexibility also allows for more generic network architectures that are more agnostic to different outputs and detector technologies.

The HEPTrkX network architecture developed at the LHC served as a proof-of-concept for GNN applications in HEP. This model considered each energy deposition in the detector as a node in a graph, and predicted binary scores on graph edges between neighbouring detector layers to perform link prediction for track-forming [11]. The continuation of this work, both within its successor collaboration, ExaTrkX [12], and elsewhere, explores more novel concepts such as hierarchical networks [13], equivariance [14] and object condensation [15].

GNNs have also been explored within neutrino physics, albeit to a lesser extent. The IceCube collaboration explored a GNN architecture for high-energy ν_μ selection, that considers each optical module in the detector as a graph node; this approach outperformed both traditional and CNN event selections [16]. Additionally, sparse CNN-based reconstruction techniques discussed previously have demonstrated the viability of utilizing GNNs at later stages to infer particle flow [17].

This paper discusses NuGraph2, a GNN developed by the ExaTrkX collaboration for low-level particle reconstruction in neutrino physics. This network architecture conceptualizes the problem of particle reconstruction similarly to its HL-LHC counterpart, but arrives at a substantially different solution due to the domain differences between collider and neutrino physics.

NuGraph2 represents a substantial refinement of a proof-of-concept architecture [18], retroactively named “NuGraph1”. This earlier model was more directly inspired by the HEPTrkX architecture, performing link prediction between detector hits. This paradigm is effective at the LHC, where particle tracks are emitted radially outwards from the center of the detector, allowing tracks to be reconstructed by connecting hits on subsequent detector layers. However, in a neutrino physics context, where no such relationship between detector geometry and particle direction can be assumed, link prediction is no longer an effective method for hit clustering.

NuGraph2 operates by classifying hits on detector planes according to particle type, leveraging nexus connections between planes to utilize 3D context information. The network presented in this paper provides two outputs – a binary filter that rejects hits that are not part of the primary physics interaction, and a semantic classifier that assigns each hit a particle type. However, the core message-passing engine discussed in this paper can be further leveraged for a broad range of applications such as event classification, clustering and more.

II. EXPERIMENTAL SETUP

The results presented in this paper utilize an open dataset released by the MicroBooNE collaboration [19]. This dataset consists of simulated ν_e and ν_μ interactions in a realistic simulation of the MicroBooNE detector, with real data hits from cosmic interactions overlaid.

MicroBooNE is a LArTPC detector operating at Fermilab [20], consisting of a cryostat containing cryogenically cooled liquid argon. Charged particles interacting in the argon volume produce ionization electrons that drift under an electric field to a series of three wire planes; these electrons induce electronic signals on the first two wire planes before being collected on the third. The three wire planes (U, V and Y) are oriented at an angle of +60 deg, -60 deg and 0 deg with respect to vertical, respectively. If the t_0 of a particle interaction is known, then the time at which signal pulses on TPC wires occur can be used as a measure of the signal’s distance from the wires. Due to the differing wire angles, each TPC wire plane provides a unique 2D representation of a particle interaction; these three 2D representations can be combined to provide a full 3D reconstruction of a particle interaction.

In order to train the NuGraph2 architecture, this dataset is used to construct heterogeneous graph objects representing each neutrino interaction; for this paper, only hits passing the Pandora [21] neutrino slice selection are considered. Pulses on TPC wires are summarized into discrete Gaussian hits occurring on specific wires at specific times. An independent subgraph is formed for each wire plane, with the hits on a given plane used as the basis for that plane’s graph nodes. Four input features are constructed for each graph node: the hit’s wire index, time coordinate, and the integral and Root Mean Square (RMS) width of its Gaussian pulse.

Edges connecting hits within each detector plane’s subgraph are formed using a Delaunay triangulation algorithm using each node’s wire and time coordinates. Other edge formation algorithms, such as k-nearest-neighbour (kNN) and radius-based approaches, were also explored, but Delaunay triangulation was found to be both the most computationally efficient and the most performant. It generates a fully connected graph with a mixture of short-range and long-range connections, which allows it to maintain performance in the presence of spatially disconnected points due to detector effects such as dead wire regions and physics features such as disconnected electromagnetic shower segments.




FIG. 1: A diagram of the operation of LArTPC detector geometry.

In addition to planar subgraphs formed from detector hits, a “nexus” subgraph is constructed using spacepoints generated by SpacePointSolver [22], an upstream algorithm that groups together 2D hits across detector planes. This grouping of planar hits is utilized to define graph edges connecting planar nodes to nexus nodes. These nexus nodes are essentially “virtual,” in that they are defined solely by the graph edges connecting them to planar hits – they have no dedicated input features, and no internal graph edges connecting nexus nodes to each other.

The construction of graph objects from simulated events is achieved using a generalized workflow developed as part of this work. The NuML package [23] is used to construct event-level HDF5 files containing tables summarizing low-level information such as simulated particles, simulated energy deposits and observed detector hits. These data structures are designed to be generalizable, and although the work described in this paper concerns the MicroBooNE detector, this workflow has also been successfully deployed in the context of other neutrino detectors, LArTPC and otherwise.

This event HDF5 format is then efficiently preprocessed into graph objects using the PyNuML package [24]. This package leverages MPI parallelism to efficiently produce a large dataset into graph objects for model training at scale. Processed graphs are stored as individual datasets of compound data types in an HDF5 file. Using the compound data type allows us to store each graph sample in a contiguous space in the file, so that reading an entire graph can be completed in a single read call. Such design effectively minimizes the number of required I/O operations and thus reduces the overall I/O cost of the data loading phase.

PyNuML is designed to offer an efficient and flexible solution for many boilerplate ML tasks in HEP. It offers algorithms to efficiently generate ground truth labels from a tree of simulated particles, and then propagate them onto the observed detector hits, an approach which can be orders-of-magnitude more efficient than the typical “backtracking” approach taken by particle physicists. The highly-parallel event processing backbone is separated out from label generation and object construction algorithms, allowing the user to easily utilize, customize or override any aspect of the processing framework for their own needs.

III. MODEL ARCHITECTURE

The NuGraph2 network architecture consists of an encoder, an iterative message-passing engine, and a set of decoders; a high-level diagram of the network structure is provided in the right diagram of Fig. 2. Each iteration of the message-passing engine consists of two components – a planar block responsible for propagating information internally within each detector plane, which is shown in the left diagram of Fig. 2 and described in detail in Sec. III B 1, and a nexus block responsible for passing information between detector planes, which is shown in the central diagram of Fig. 2 and described in detail in Sec. III B 2.

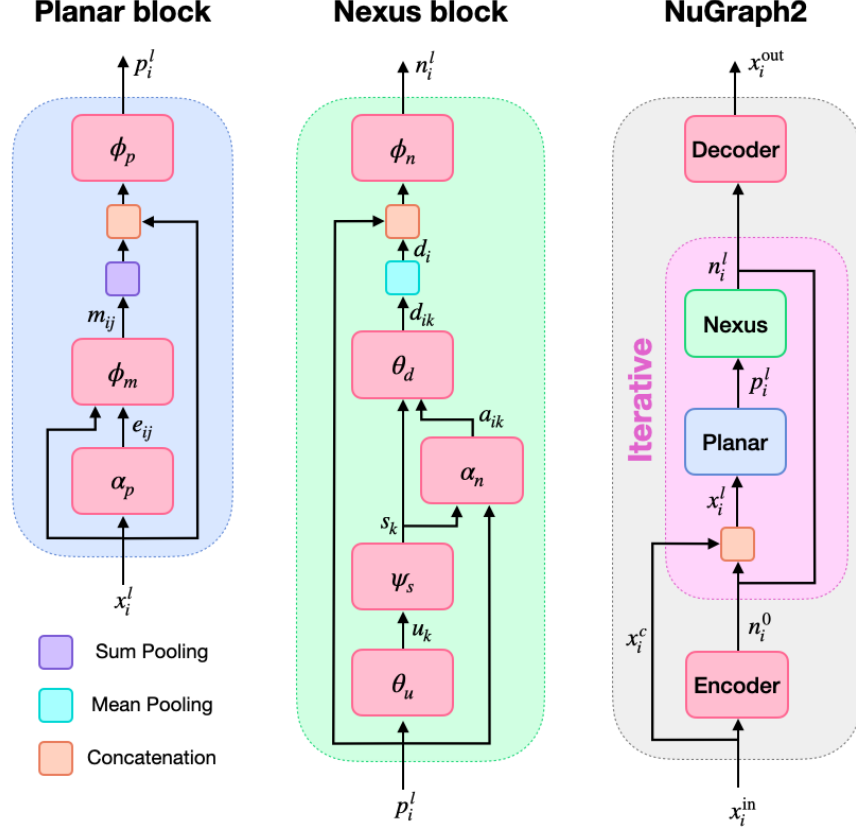


FIG. 2: Block diagram for the planar (left) and nexus (center) blocks, and the overall structure of the NuGraph2 network architecture (right). ϕ denotes a planar convolution operator, ψ a nexus convolution operator, θ a convolution that transforms between planar and nexus nodes, and α a categorical cross-attention convolution. The i subscript refers to planar nodes, $j \in \mathcal{N}(i)$ their neighbouring planar nodes, and $k \in \mathcal{X}(i)$ the nexus nodes connected to each planar node; the l superscript refers to the current iteration of message-passing.

All linear convolution layers are immediately followed by an activation function. The softmax and sigmoid functions are used for layers explicitly producing categorical and binary probabilities, respectively. For all other convolutional layers, a hyperbolic tangent activation function is used, as it was found to provide superior performance and stability when compared to the ReLU and Mish [25] activation functions.

NuGraph2 utilizes a **categorical embedding** mechanism, in which the network learns a distinct set of hidden features for each particle category, as defined by hit semantic labels. Internal graph node feature tensors have shape (N_n, N_c, N_f) , where N_n is the number of graph nodes, N_c is the number of semantic particle categories, and N_f is the number of latent features per category.

Convolution layers operating on these categorical tensors utilize a **categorical linear** mechanism, in which a separate linear layer operates on the hidden features for each category, in a manner equivalent to grouped linear convolutions [26]. A categorical linear convolution layer requires a factor of N_c fewer neurons than an equivalent fully connected linear layer.

The result of this categorical convolution mechanism is that each category’s hidden features are effectively “unaware” of the corresponding features for other particle categories, meaning that the network can learn information specific to each particle class.

We also utilize a specialized multi-head self-attention mechanism referred to as **categorical cross-attention**. This mechanism applies an MLP to a node feature tensor in a categorical embedding, convolving down to a single feature for each node category, and then applying a softmax across the categorical dimension. This produces a set of probabilities per category, which can then be applied to the original node feature tensor as self-attention weights.

When combined, the categorical linear and categorical cross-attention mechanisms provide an efficient and effective convolution mechanism. The application of categorical linear convolutions means that each category’s features are

allowed to evolve independently, while the categorical cross-attention mechanism allows for some controlled exchange of context information between categories. Categorical cross-attention is utilized to weight messages formed on graph edges during message-passing, suppressing information from categories that are not strongly activated while enhancing information from categories with strong activation.

Throughout this section, different types of convolution operator are denoted with different characters: ϕ for 2D planar convolutions, ψ for 3D nexus convolutions, θ for convolutions that transform between planar and nexus nodes, and α for categorical cross-attention convolutions.

A. Encoder

NuGraph2’s encoder serves to take the input graph node features x^{in} and generate the initial categorical node embedding. First, the (N_n, N_f) input tensor is simply repeated N_c times across a new second dimension to produce a new tensor x^c with shape (N_n, N_c, N_f) . This is then used to generate the initial node embedding

$$n^0 = \phi_e(x^c), \quad (1)$$

where ϕ_e is a categorical linear multi-layer perceptron (MLP).

B. Message passing

After initial encoding to produce graph node features in a categorical embedding space, the network utilizes an iterative two-step message-passing mechanism as its core engine. This engine consist of a set of **planar blocks**, that operate independently on each 2D detector plane’s subgraph, and a **nexus block** that exchanges information between 2D planes to allow for degeneracy breaking.

At the start of each message-passing iteration, graph node features at iteration l , n_i^l , are concatenated with x^c , the input graph node features prior to any convolutions repeated across the categorical dimension, to form the x_i^l tensor.

1. Planar block

The **planar block** is responsible for propagating information internally inside each detector plane. The input features on each plane x^l in a categorical embedding are convolved to produce a set of categorical cross-attention scores

$$e_{ij} = \alpha_p(x_i^l, x_j^l), \quad (2)$$

where α_p is a categorical cross-attention block, and $j \in \mathcal{N}(i)$ are the neighbors of node i . The resulting attention weights are then used to weight messages formed on graph edges,

$$m_{ij} = \phi_m(x_i, x_j, e_{ij}), \quad (3)$$

where ϕ_m is a categorical linear MLP. These messages are then aggregated into graph nodes via sum pooling,

$$m_i = \sum_{j \in \mathcal{N}(i)} m_{ij}, \quad (4)$$

and then concatenated with the input node features and convolved once more,

$$p_i^l = \phi_p(m_i^l, x_i^l), \quad (5)$$

where ϕ_p is a categorical linear MLP and p_i^l are updated planar node features.

2. Nexus block

Following planar convolutions, the **nexus block** is responsible for mixing information between detector planes in order to break degeneracies inherent in each 2D representation. As an example, a long minimum ionising track (ie.

a muon) propagating at a shallow angle with respect to the pitch of the wires on one detector plane can appear as a shorter, more highly ionising track (ie. a proton). The nexus block allows the network to cross-reference this representation with representations in other planes.

As discussed in Sec. II, nexus nodes are “virtual” in the sense that they are defined solely by their connections to planar nodes. No message-passing is performed internally in the nexus space; instead, planar node features are propagated into the nexus nodes via message-passing, and the resulting nexus node features are convolved to form new nexus node features containing a mixture of information from all planes. These nexus features are then propagated back down to the planar nodes via message-passing, providing each planar subgraph with an injection of contextual information from other planes.

In order to accomplish this, node features in each planar subgraph are first used to form messages on the edges connecting the planar and nexus nodes

$$u_k = \theta_u(p_i^l), \quad (6)$$

where $k = \mathcal{X}(i)$ are the nexus nodes connected to planar node i , and θ_u is a simple message-passing convolution operator that transforms planar node features to nexus nodes. By construction each nexus node is connected to a maximum of one node on each plane, and so planar features are aggregated into nexus nodes via concatenation. In the event that a nexus node is not connected to any node on a given plane, that plane’s contribution to the nexus node’s features are set to zero.

Nexus node features are then convolved to produce new nexus features

$$s_k = \psi_s(u_k), \quad (7)$$

where ψ_s is a categorical linear MLP. Nexus node features are then passed back down to each planar subgraph, first by producing categorical cross-attention weights

$$a_{ik} = \alpha_n(p_i^l, s_k), \quad (8)$$

where α_n is a categorical cross-attention MLP, and then by using those weights to pass messages formed on edge features

$$d_{ik} = \theta_d(s_k, a_{ik}), \quad (9)$$

where θ_d is a categorical linear MLP. Messages are then aggregated into planar nodes using mean pooling,

$$d_i = \frac{1}{|\mathcal{X}(i)|} \sum_{k \in \mathcal{X}(i)} d_{ik}. \quad (10)$$

Here mean pooling is utilized instead of the sum pooling used in Eqn. 4 for planar message-passing, as it yields stronger performance. These new planar node features are then shortcut connected with previous node features and convolved,

$$n_i^l = \phi_n(p_i^l, d_i^l), \quad (11)$$

where ϕ_l is a categorical linear MLP.

C. Decoders

The output of the final message-passing iteration is passed into two decoders, which are responsible for producing the filter and semantic outputs. The filter decoder first flattens the categorical embedding tensor from a three-dimensional tensor with shape (N_n, N_c, N_f) to a two-dimensional tensor with shape $(N_n, N_c \times N_f)$, and then convolves each nodes’s features into a single binary classifier score by utilising a single linear convolutional layer followed by a sigmoid activation. The semantic decoder is similar in structure but retains the categorical embedding, convolving each category’s features for each node down to a single score, and then taking the softmax activation across categories to produce categorical probability scores for each node.

IV. TRAINING

NuGraph2 utilizes PyTorch [27] with PyG [28] for our model, leveraging the model training infrastructure provided by PyTorch Lightning. Training is optimized using the AdamW optimizer [29] with a learning rate of 0.001, scheduled using the OneCycleLR scheduler [30] for a total of 80 epochs.

A simple binary cross-entropy is used to generate the loss for the hit filtering decoder. For the semantic decoder, we utilize the Recall Loss [31], as it provides an optimal balance of precision and recall performance even in the presence of severe class imbalance. This approach was found to provide significantly more balanced performance than other techniques explored, such as the inverse frequency-weighted Categorical Cross-Entropy Loss, which naively optimized recall for minority classes at the expense of precision, and the Focal Loss [32], which failed to learn on minority classes altogether.

Loss functions for multi-task training are balanced by utilising a temperature mechanism to automatically find an optimal performance balance between both tasks [33].

V. RESULTS

A. Model performance

Network performance is quantified in terms of **recall** (efficiency) and **precision** (purity), which are defined as

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \qquad \text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \qquad (12)$$

where TP, FP and FN represent true positives, false positives and false negatives, respectively. We also utilize the **receiver-operator characteristic** (ROC) curve, in which the true positive rate and false positive rate are computed and visualized for a range of probability score thresholds.

For the network’s binary filter output, positive and negative predictions are determined per hit by applying a probability threshold of 0.5. The filter decoder achieves a recall of 0.980 and a precision of 0.979; a recall confusion matrix is shown in Fig. 3. Due to the preponderance of signal hits over background hits in the training set, the recall metric for signal hits is larger than for background hits when applying a balanced threshold of 0.5. The filter output can be optimized towards selecting signal hits or rejecting background hits by applying a different threshold, as reflected in the filter ROC curve shown in Fig. 6a.

For the network’s semantic probability outputs, a discrete class prediction per hit is determined by selecting the class with the largest probability score. The semantic decoder achieves an overall recall and precision of 0.949; these values are equal due to the scheme used to weight the loss function. Recall and precision confusion matrices are shown in Fig. 4.

The semantic predictions are most performant for the MIP class, which is the most highly-represented class in the training dataset. This is due to the abundance of muons in the training set, coupled with the fact that muons are minimum-ionizing particles that produce long tracks, and therefore many detector hits. Conversely, the semantic predictions are least accurate for the Michel class, which is the least represented class in the dataset, due to a relative scarcity and low hit multiplicity.

The most common modes of semantic misclassification are consistent with qualitative expectations. The HIP and MIP classes are topologically similar, as they both form particle tracks, and they are typically distinguished through differences in track length and energy deposition. One of the network’s primary failure modes is the misclassification of HIP hits as MIP-like – 4.7% of true HIP hits are predicted to be MIP-like.

The shower, Michel and diffuse classes have similar conceptual overlap, as all three more frequently consist of smaller energy depositions – electromagnetic showers are typically formed from a cascade of smaller and frequently disconnected energy depositions, while Michel electrons are short tracks, and diffuse activity often manifests as small and disconnected energy depositions due to activity from neutrons and EM activity such as Compton scatters. This is reflected in the network output, as 11% and 10% of true Michel hits are misclassified as shower-like and diffuse respectively. Similarly, 5.8% and 2.1% of true diffuse hits are misclassified as shower-like and Michel-like, respectively.

Another primary failure mode is between the MIP and Michel classes, which are strongly correlated as Michel electrons always occur at the end of a muon track. This confusion primarily manifests as true MIP hits being misclassified as Michel-like, which is a consequence of the weighting scheme applied to the loss function. True MIP hits are misclassified as Michel-like at a rate of only 0.65%, but due to class imbalance in the dataset, these misclassified hits constitute 26% of hits predicted to be Michel-like by the network. By contrast, only 6% of Michel hits are misclassified as MIP-like, and these hits only constitute 0.1% of predicted MIP hits.




FIG. 3: Confusion matrix for hit filtering, with probabilities normalized per predicted class in order to highlight recall.

Predicted score distributions for each semantic class are shown in Fig. 5, while ROC curves for each class are shown in Fig. 6b.

For semantic hit labelling, we also define a metric **consistency** as the proportion of nexus nodes for which all connected planar nodes share the same semantic label, in order to quantify the compatibility of semantic classification across different detector planes. Semantic labels in ground truth have a consistency of 98.04%, while semantic labels predicted by NuGraph2 have a consistency of 94.81%. By contrast, a planar-only model that considers each detector plane as an independent graph by omitting the nexus phase of message-passing achieves a consistency of only 67.38%, demonstrating the nexus block’s role in producing a coherent representation across detector planes.


B. Inference benchmarks

Inference benchmarks were carried out on an Nvidia DGX GPU cluster. Evaluating NuGraph2 model output event-by-event on CPU takes 0.12 s/event. Inference on GPU reaches 0.005 s/event for sufficiently large batch sizes; GPU inference time as a function of batch size is shown in Fig. 7.


NuGraph2 inference has also been implemented in the LArSoft software framework [34], using a TorchScript-compiled model running on CPU. This preliminary implementation achieves an inference time of 0.3 s/event, which accounts for both graph construction and model inference time. Peak CPU memory utilization is currently 3.1 GB; memory overhead can be further improved, but this work is beyond the scope of this paper. A mature and flexible inference workflow that utilizes the SONIC interface [35] to run graph construction and model inference natively in Python is currently in development.

C. Hyperparameter optimization

The number of hidden parameters in each NuGraph2 convolution layer are configurable via hyperparameters, implemented independently for planar convolutions and nexus convolutions. The default values were chosen for a balance of performance and memory overhead. Network performance as a function of the number of node features and nexus features is shown in Fig. 8. Network performance is primarily dependent on the size of the planar embedding, with a weaker dependence on the size of the nexus embedding. The nominal sizes of 64 and 16 features for the planar



(a) Confusion matrix normalized by true semantic class.



(b) Confusion matrix normalized by predicted semantic class.

FIG. 4: Confusion matrices for hit semantic segmentation.




FIG. 5: Predicted score distributions for each semantic category, separated into detector hits that belong to that category in truth (left), and hits that belong to other categories in truth (right). These distributions are area-normalized so their shapes can be compared.




FIG. 6: Receiver-Operator Characteristic (ROC) curves for model predictions.

and nexus embeddings, respectively, were chosen as larger embeddings incur significant memory overhead costs while network performance plateaus. More details on all model variants are provided in Tab. I.

D. Event displays

In order to provide context for the model performance metrics described above, this section discusses NuGraph2’s performance on a series of representative events, designed to highlight events where the network performs well, and also events that highlight common failure modes. All events shown are drawn from the test partition of the dataset, and so were not used during network training.

| Variant | Planar feats. | Nexus feats. | Filter recall | Semantic recall | Train GPU mem. (GB) |
|------------|---------------|--------------|---------------|-----------------|---------------------|
| nominal | 64 | 16 | 0.978 | 0.947 | 6.19 |
| planar-4 | 4 | 16 | 0.963 | 0.879 | 0.69 |
| planar-8 | 8 | 16 | 0.968 | 0.899 | 1.12 |
| planar-16 | 16 | 16 | 0.971 | 0.919 | 1.83 |
| planar-32 | 32 | 16 | 0.975 | 0.935 | 3.29 |
| planar-128 | 128 | 16 | 0.979 | 0.950 | 11.98 |
| nexus-4 | 64 | 4 | 0.978 | 0.943 | 6.18 |
| nexus-8 | 64 | 8 | 0.977 | 0.944 | 6.19 |
| nexus-32 | 64 | 32 | 0.978 | 0.946 | 6.19 |

TABLE I: Network performance and efficiency metrics for different model configurations.




FIG. 7: Inference time per graph on GPU as a function of batch size.




FIG. 8: Network performance as a function of number of hidden features utilized.

For the purposes of this section, the network’s filter prediction is converted into a binary score by applying a threshold at 0.5. This threshold is configurable, and can be varied to optimize towards signal selection or background rejection. Similarly, the network’s classwise semantic probabilities are reduced to a discrete semantic label for each hit by selecting the class with the highest probability.

1. Muon interaction

The first representative event is Run 7049, Subrun 391, Event 19599 of the MicroBooNE open data release, shown in Fig. 9. This event is a charged current quasielastic (CCQE) ν_μ interaction, consisting of a longer track produced by a muon with 378 MeV momentum, and a shorter track produced by a proton with 740 MeV momentum. This event was selected as an example of a topologically simple event that is well-reconstructed by the network.

This event consists of 1058 hits across all wire planes, of which 777 are hits from the simulated neutrino interaction. The filter decoder correctly classifies 1053 hits; of the 5 misclassified hits, 2 are signal hits misclassified as background, and 3 are background hits misclassified as signal. The semantic decoder misclassifies 6 of the 777 signal hits – 5 MIP hits are misclassified (2 as HIP, 2 as Michel and 1 as diffuse), and a single HIP hit is misclassified as diffuse.

This neutrino interaction is a clean two-prong topology, and the network is able to reject background hits and classify signal hits with $> 99\%$ efficiency.

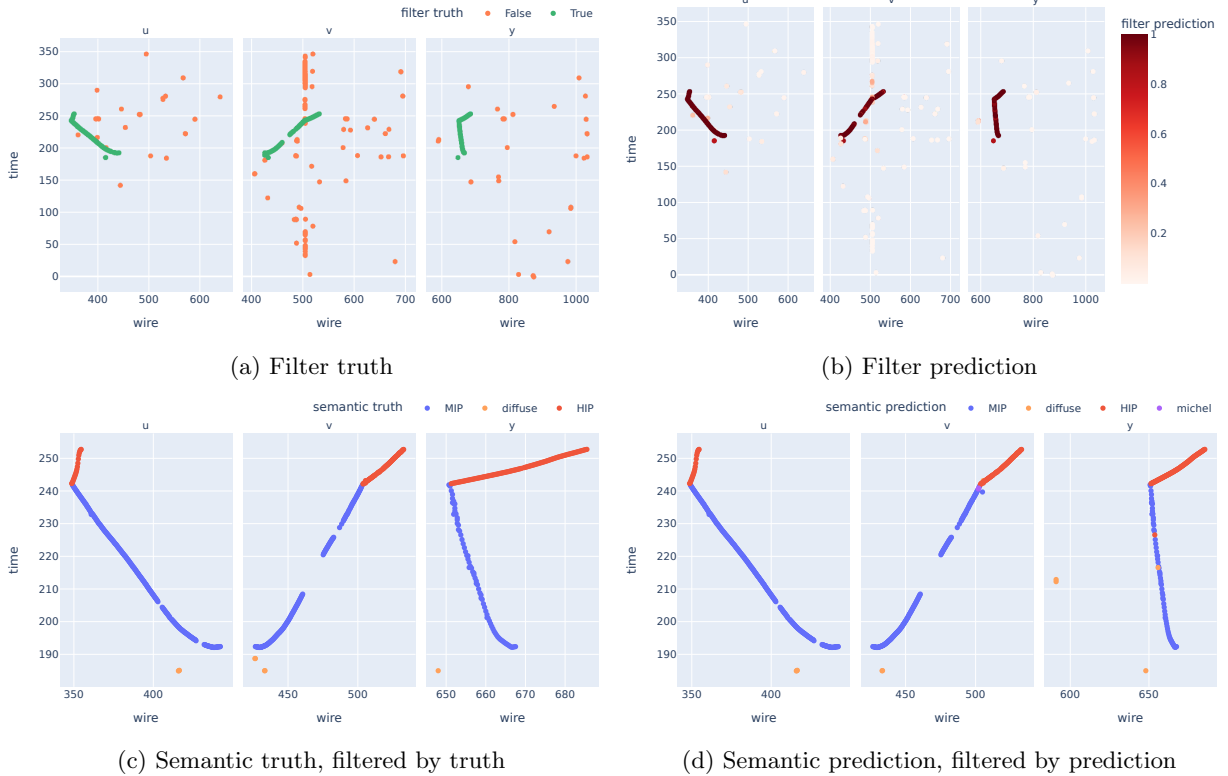


FIG. 9: Event displays for the first representative event: Run 7049, Subrun 391, Event 19599 of the MicroBooNE open data release.

2. Shower interaction

The second representative event is Run 6999, Subrun 11, Event 595 of the MicroBooNE open data release, shown in Fig. 10. This event is a neutral current (NC) interaction, consisting of a track produced by a proton with 719 MeV momentum and a π^0 meson with 91 MeV momentum that is not observed directly, but decays into a pair of photons with momenta of 60 MeV and 103 MeV, each of which produces an electromagnetic (EM) shower. This event was selected as an example of a topologically simple event containing EM showers that is well-reconstructed by the network.

This event consists of 568 hits across all wire planes, of which 469 are hits from the simulated neutrino interaction. The filter decoder correctly classifies 566 hits, and misclassifies 2 background hits as signal. The semantic decoder correctly classifies all signal hits.

This event is easy to reconstruct, as background hits are largely disconnected from the physics interaction, and the photon showers are offset from the neutrino vertex. Under these conditions, the network is able to achieve near-perfect performance.

3. Michel electron

The third representative event is Run 6110, Subrun 152, Event 7630 of the MicroBooNE open data release, shown in Fig. 11. This event is a CCQE ν_μ interaction consisting of a single track produced by a μ^+ with 542 MeV momentum, that decays into a Michel electron with 29 MeV momentum. This event was selected as an example of a well-reconstructed ν_μ interaction containing a Michel electron.

This event consists of 739 hits across all wire planes, of which 711 are hits from the simulated neutrino interaction. The filter decoder correctly classifies 738 hits, and misclassifies 1 signal hit as background. The semantic decoder misclassifies 3 of the 711 signal hits – 2 MIP hits are misclassified as Michel, and a single Michel hit is misclassified as MIP.

This interaction is a clean Michel electron topology, and the network is able to reject background hits and classify signal hits with $> 99\%$ efficiency.

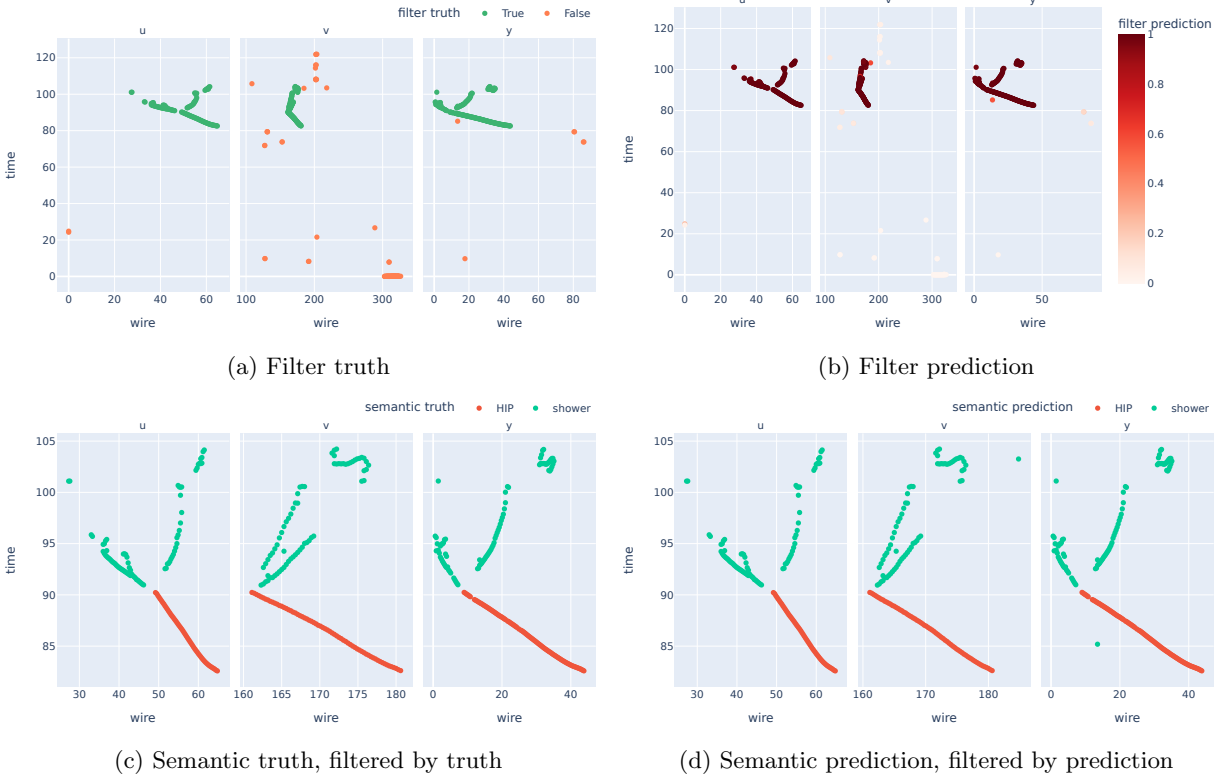


FIG. 10: Event displays for the second representative event: Run 6999, Subrun 11, Event 595 of the MicroBooNE open data release.

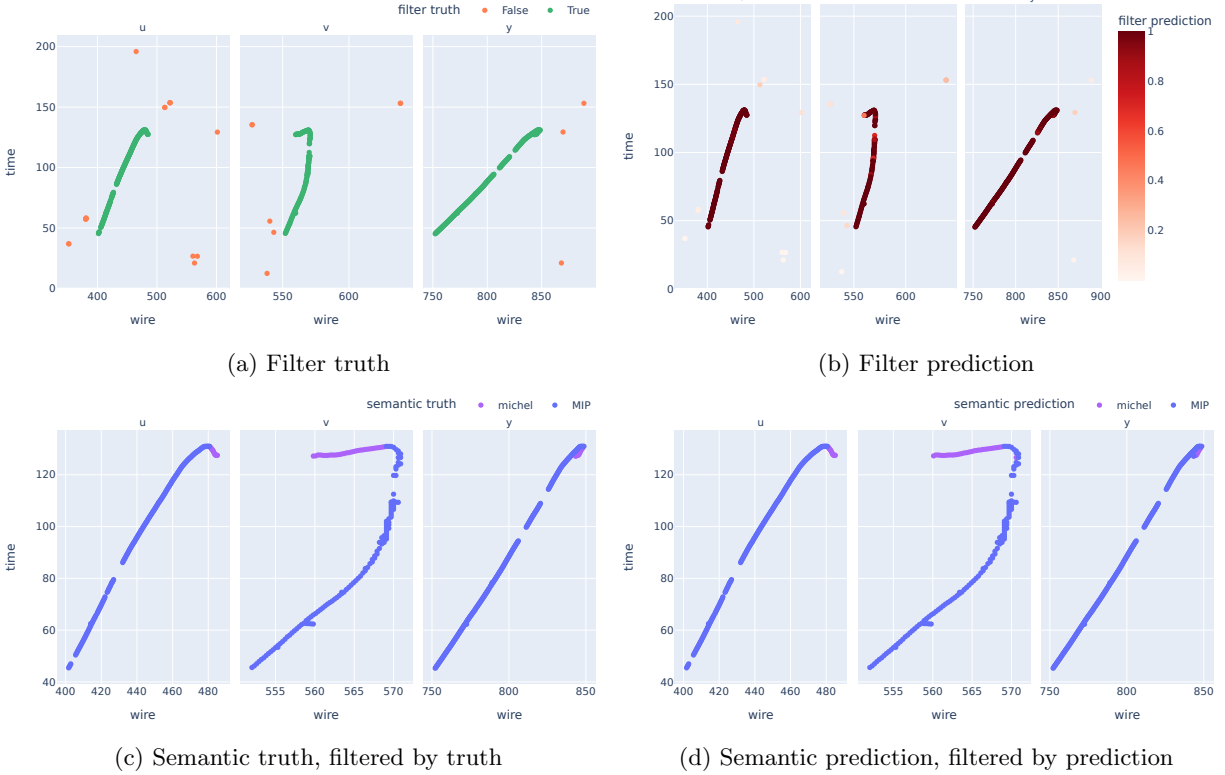


FIG. 11: Event displays for the third representative event: Run 6110, Subrun 152, Event 7630 of the MicroBooNE open data release.

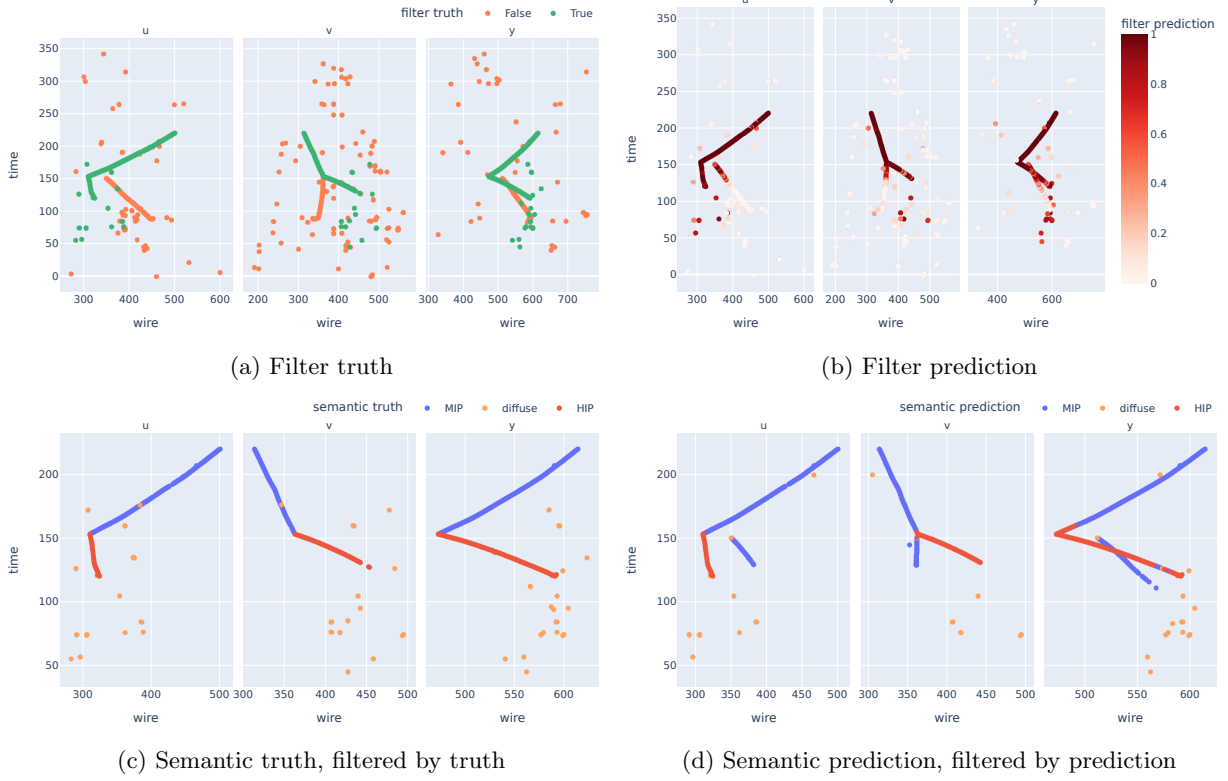


FIG. 12: Event displays for the fourth representative event: Run 5459, Subrun 94, Event 4738 of the MicroBooNE open data release.

4. Intersecting cosmic

The fourth representative event is Run 5459, Subrun 94, Event 4738 of the MicroBooNE open data release, shown in Fig. 12. This event is a CCQE ν_μ interaction consisting of a longer track produced by a muon with 944 MeV momentum and a shorter track produced by a proton with 1.29 GeV momentum; an unrelated cosmic track also intersects the neutrino event.

This event consists of 3138 hits across all wire planes, of which 2069 are hits from the simulated neutrino interaction. The filter decoder correctly classifies 2771 hits; of the 367 misclassified hits, 118 are signal hits misclassified as background and 249 are background hits misclassified as signal. The semantic decoder misclassifies 133 of the 2069 signal hits – 36 MIP hits are misclassified as HIP, 83 HIP hits are misclassified (74 as HIP and 9 as diffuse), and 14 diffuse hits are misclassified (13 as MIP and 1 as HIP).

The majority of background hits misclassified as signal by the model belong to the cosmic track that intersects the physics interaction. The proximity of this cosmic track to the simulated neutrino interaction makes it difficult to identify as background, as it points close to the neutrino vertex in 3D. Despite this, the network is able to correctly identify the majority of the track’s hits as background, and only misidentifies the track segment closest to the neutrino interaction. The semantic decoder also misclassifies hits at the start of the muon track as being more proton-like, especially in the V and Y planes. It is also worth noting that the network’s semantic prediction correctly identifies the background hits from the cosmic track as MIP-like, indicating that the semantic information learned on signal hits from simulation is transferable to background hits from real data.

5. High-multiplicity event

The fifth representative event is Run 6780, Subrun 200, Event 10006 of the MicroBooNE open data release, shown in Fig. 13. This event is a CC ν_μ interaction consisting of a μ^- with momentum 134 MeV, 5 protons with momenta 777 MeV, 145 MeV, 99 MeV, 318 MeV and 177 MeV, two π^+ with momenta 434 MeV and 191 MeV, and two π^0 with momenta 284 MeV and 113 MeV, which are not directly observed but decay into photon pairs with momenta 84 MeV

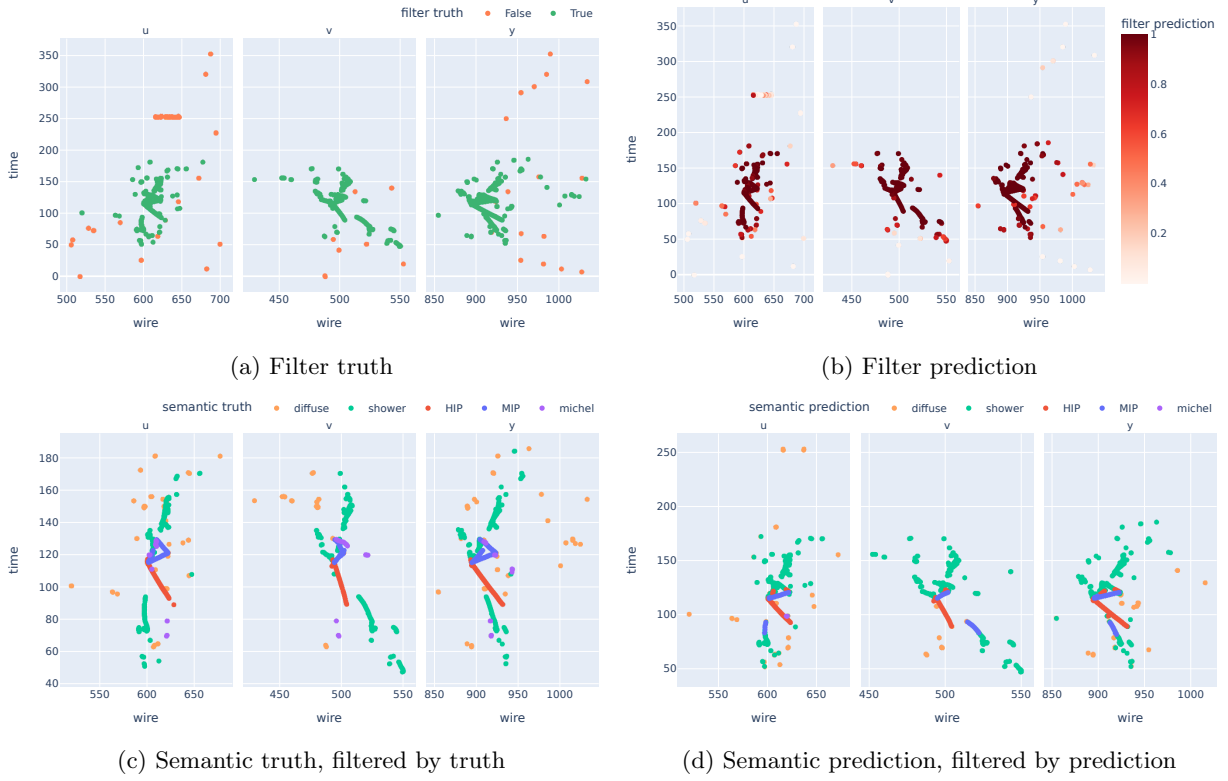


FIG. 13: Event displays for the fifth representative event: Run 6780, Subrun 200, Event 10006 of the MicroBooNE open data release.

and 231 MeV, and 32 MeV and 144 MeV, respectively. This event was selected as an example of a high-multiplicity event containing hits from all five semantic classes, that is more challenging for the network to reconstruct.

This event consists of 1807 hits across all wire planes, of which 1657 hits from the simulated neutrino interaction. The filter decoder correctly classifies 1765 hits; of the 42 misclassified hits, 24 are signal hits misclassified as background, and 18 are background hits misclassified as signal. The semantic decoder misclassifies 477 of the 1657 signal hits – 196 MIP hits (60 as HIP and 136 as shower), 15 HIP hits (14 as shower and 1 as diffuse), 73 shower hits (56 as MIP, 1 as HIP and 16 as diffuse), 88 Michel hits (1 as HIP, 67 as shower and 20 as diffuse) and 105 diffuse hits (2 as HIP, 99 as shower and 4 as Michel).

Events with a higher particle multiplicity are more challenging to reconstruct, as detector hits produced by different particles are more likely to overlap. Additionally, the MicroBooNE open dataset primarily consists of neutrino interactions with a low particle multiplicity, and events with ≥ 10 primary particles constitute only 10% of the training set. Considering these limitations, the network performs respectably in correctly semantically labelling over two thirds of the hits in this event.

The 777 MeV momentum proton leaves a long track, which is correctly classified by the model, as is the shorter track produced by a secondary proton with 341 MeV momentum. The primary proton with 318 MeV produces a short track (only 13 hits over all wire planes) and is misclassified by the model as shower-like.

The network performs most poorly on the MIP class. The primary muon track is entirely misclassified, with its hits being assigned to a combination of the shower and HIP classes. The π^+ with 434 MeV momentum is correctly identified as MIP-like; however, it then inelastically scatters to form a new π^+ , which decays into a μ^+ and finally a Michel positron. These secondary tracks are all misclassified by the model as shower-like, as they overlap with the EM showers produced by the decay of π^0 primaries.

The model identifies the majority of hits belonging to all EM showers, with the exception of one shower root, which it classifies as MIP-like. This is technically correct from a physics perspective, as the showering photon pair-produces into an electron-positron pair that initially appear track-like before undergoing bremsstrahlung and initiating an electromagnetic cascade, but it is not consistent with the semantic truth labelling scheme as currently defined. Such failure modes can likely be resolved in future by redesigning the ground truth labelling scheme to account for them.

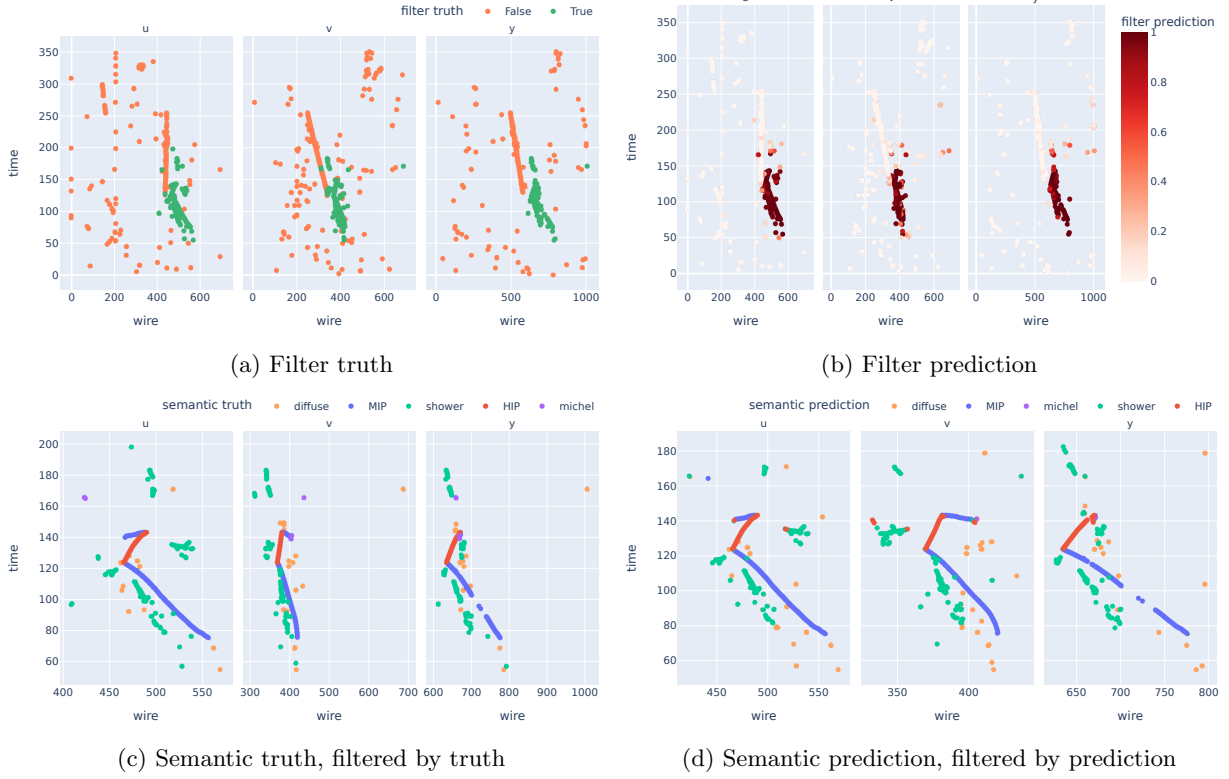


FIG. 14: Event displays for the sixth representative event: Run 6343, Subrun 72, Event 3606 of the MicroBooNE open data release.

6. High-multiplicity event

The sixth representative event is Run 6343, Subrun 72, Event 3606 of the MicroBooNE open data release, shown in Fig. 14. This event is a CC ν_μ interaction that produces a μ^- track with 456 MeV momentum, a Λ baryon with 701 MeV momentum and a K^+ meson with 497 MeV momentum. The Λ produces neutron activity, a proton track with 169 MeV momentum, and a π^0 with 180 MeV momentum that decays into a photon pair with momenta 178 MeV and 47 MeV. The K^+ produces a π^0 with 205 MeV momentum that decays into a photon pair with momenta 156 MeV and 90 MeV, and also a π^+ with 205 MeV momentum that decays into a μ^+ and then a Michel e^+ . This event was selected as a challenging high-multiplicity interaction on which NuGraph2 demonstrates good performance.

This event consists of 3682 hits across all wire planes, of which 2036 are hits from the simulated neutrino interaction. The filter decoder correctly identifies 3527 hits; of the 155 misclassified hits, 124 are signal hits misclassified as background, and 31 are background hits misclassified as signal. The semantic decoder misclassifies 107 of the 2036 signal hits – 18 MIP hits are misclassified (5 as HIP, 10 as shower, 1 as Michel and 2 as diffuse), 3 HIP hits are misclassified as shower, 53 shower hits are misclassified (3 as MIP, 8 as HIP, 6 as Michel and 36 as diffuse), 15 Michel hits are misclassified (1 MIP, 1 as HIP, 3 as shower and 10 as diffuse) and 18 diffuse hits are misclassified (17 as shower, 1 as Michel).

The network is able to identify all visible particles in the event, although not all of them are perfectly reconstructed. In regions where energy deposition from EM showers are more sparse, the network exhibits confusion between the shower and diffuse categories. The model correctly identifies the hits associated with the K^+ track, and the majority of the subsequent π^+ track, but misclassifies some hits in the rapid decay to μ^+ and then e^+ , with some hits being identified as shower-like or diffuse. As with the previous example event, these failure modes can potentially be significantly improved upon by expanding the network to explicitly generate particle clustering as a learning target.

VI. SUMMARY

The NuGraph2 architecture is a state-of-the-art graph neural network for particle reconstruction in neutrino physics, achieving 98.0% efficiency in filtering background hits and 94.9% efficiency in semantically labelling hits, while producing 2D representations that are 94.8% consistent. Unlike CNN-based techniques, it is capable of operating on the detector observables in their native structure, without requiring any arbitrary pixelization, voxelization, truncation or downsizing.

The network is also designed to be maximally agnostic to detector technology. It is able to operate on any data structure consisting of detector hits, and offers flexibility in terms of the number of detector planes, truth labelling schemes, and number of graph node input features, rendering it simple to apply to other detector geometries, both LArTPC and otherwise.

NuGraph2's core message-passing engine also provides a general-purpose solution for a variety of tasks. The two decoders presented in this paper are implemented in a modular manner, meaning the architecture can easily be extended or adapted to leverage the core engine for additional learning tasks. Several additional applications are already in active development, utilizing hierarchical graph convolution methods to reconstruct particle and event level quantities, and this work will be discussed in future publications.

ACKNOWLEDGMENTS

This research was supported in part by the U.S. Department of Energy's Office of Science, Office of High Energy Physics, under Contract No. DE-AC02-07CH11359 (CompHEP Exa.TrkX). We would also like to acknowledge the help and support of students working as part of The University of Chicago's Data Science Institute Clinic and Summer Lab programs. We acknowledge the MicroBooNE Collaboration for making publicly available the data sets [36, 37] 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 [20].

-
- [1] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle, A convolutional neural network neutrino event classifier, *Journal of Instrumentation* **11** (09), P09001.
 - [2] F. Psihas, E. Niner, M. Groh, R. Murphy, A. Aurisano, A. Himmel, K. Lang, M. D. Messier, A. Radovic, and A. Sousa, Context-Enriched Identification of Particles with a Convolutional Network for Neutrino Events, *Phys. Rev. D* **100**, 073005 (2019), arXiv:1906.00713 [physics.ins-det].
 - [3] B. Abi *et al.* (DUNE), Neutrino interaction classification with a convolutional neural network in the DUNE far detector, *Phys. Rev. D* **102**, 092003 (2020), arXiv:2006.15052 [physics.ins-det].
 - [4] A. Abed Abud *et al.* (DUNE), Separation of track- and shower-like energy deposits in ProtoDUNE-SP using a convolutional neural network, *Eur. Phys. J. C* **82**, 903 (2022), arXiv:2203.17053 [physics.ins-det].
 - [5] C. Adams *et al.* (MicroBooNE), Deep neural network for pixel-level electromagnetic particle identification in the MicroBooNE liquid argon time projection chamber, *Phys. Rev. D* **99**, 092001 (2019), arXiv:1808.07269 [hep-ex].
 - [6] P. Abratenko *et al.* (MicroBooNE), Convolutional neural network for multiple particle identification in the MicroBooNE liquid argon time projection chamber, *Phys. Rev. D* **103**, 092003 (2021), arXiv:2010.08653 [hep-ex].
 - [7] C. Choy, J. Gwak, and S. Savarese, 4d spatio-temporal convnets: Minkowski convolutional neural networks, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019) pp. 3075–3084.
 - [8] D. H. Koh, P. C. de Soux, L. Dominé, F. Drielsma, R. Itay, Q. Lin, K. Terao, K. V. Tsang, and T. Usher, Scalable, proposal-free instance segmentation network for 3d pixel clustering and particle trajectory reconstruction in liquid argon time projection chambers (2020), arXiv:2007.03083 [physics.ins-det].
 - [9] P. Abratenko *et al.* (MicroBooNE), Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE, *Phys. Rev. D* **103**, 052012 (2021), arXiv:2012.08513 [physics.ins-det].
 - [10] A. Abed Abud (DUNE), Sparse Convolutional Neural Networks for particle classification in ProtoDUNE-SP events, *J. Phys. Conf. Ser.* **2438**, 012125 (2023).
 - [11] S. Farrell *et al.*, Novel deep learning methods for track reconstruction, in *4th International Workshop Connecting The Dots 2018* (2018) arXiv:1810.06111 [hep-ex].
 - [12] X. Ju *et al.* (Exa.TrkX), Performance of a geometric deep learning pipeline for HL-LHC particle tracking, *Eur. Phys. J. C* **81**, 876 (2021), arXiv:2103.06995 [physics.data-an].
 - [13] R. Liu, P. Calafiura, S. Farrell, X. Ju, D. T. Murnane, and T. M. Pham, Hierarchical Graph Neural Networks for Particle Track Reconstruction, in *21th International Workshop on Advanced Computing and Analysis Techniques in Physics Research: AI meets Reality* (2023) arXiv:2303.01640 [hep-ex].
 - [14] D. Murnane, S. Thais, and A. Thete, Equivariant Graph Neural Networks for Charged Particle Tracking, in *21th In-*

- ternational Workshop on Advanced Computing and Analysis Techniques in Physics Research: AI meets Reality (2023)* arXiv:2304.05293 [physics.ins-det].
- [15] K. Lieret and G. DeZoort, An Object Condensation Pipeline for Charged Particle Tracking at the High Luminosity LHC, in *26th International Conference on Computing in High Energy & Nuclear Physics (2023)* arXiv:2309.16754 [physics.data-an].
 - [16] N. Choma, F. Monti, L. Gerhardt, T. Palczewski, Z. Ronaghi, P. Prabhat, W. Bhimji, M. M. Bronstein, S. R. Klein, and J. Bruna, Graph neural networks for icecube signal classification, in *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA) (2018)* pp. 386–391.
 - [17] F. Drielsma, K. Terao, L. Dominé, and D. H. Koh, Scalable, End-to-End, Deep-Learning-Based Data Reconstruction Chain for Particle Imaging Detectors, in *34th Conference on Neural Information Processing Systems (2021)* arXiv:2102.01033 [hep-ex].
 - [18] V. Hewes *et al.* (Exa.TrkX), Graph neural network for object reconstruction in liquid argon time projection chambers, EPJ Web of Conferences **251**, 03054 (2021).
 - [19] G. Cerati (MicroBooNE), MicroBooNE Public Data Sets: a Collaborative Tool for LArTPC Software Development, in *26th International Conference on Computing in High Energy & Nuclear Physics (2023)* arXiv:2309.15362 [hep-ex].
 - [20] R. Acciarri *et al.* (MicroBooNE), Design and Construction of the MicroBooNE Detector, JINST **12** (02), P02017, arXiv:1612.05824 [physics.ins-det].
 - [21] R. Acciarri *et al.* (MicroBooNE), The Pandora multi-algorithm approach to automated pattern recognition of cosmic-ray muon and neutrino events in the MicroBooNE detector, Eur. Phys. J. C **78**, 82 (2018), arXiv:1708.03135 [hep-ex].
 - [22] B. Abi *et al.* (DUNE), Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics (2020), arXiv:2002.03005 [hep-ex].
 - [23] <https://github.com/vhewes/numl>.
 - [24] <https://github.com/vhewes/pynuml>.
 - [25] D. Misra, Mish: A self regularized non-monotonic activation function (2020), arXiv:1908.08681 [cs.LG].
 - [26] A. Krizhevsky, I. Sutskever, and G. E. Hinton, Imagenet classification with deep convolutional neural networks, Commun. ACM **60**, 84–90 (2017).
 - [27] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, Pytorch: An imperative style, high-performance deep learning library (2019), arXiv:1912.01703 [cs.LG].
 - [28] M. Fey and J. E. Lenssen, Fast graph representation learning with pytorch geometric (2019), arXiv:1903.02428 [cs.LG].
 - [29] I. Loshchilov and F. Hutter, Decoupled weight decay regularization (2019), arXiv:1711.05101 [cs.LG].
 - [30] L. N. Smith and N. Topin, Super-convergence: Very fast training of neural networks using large learning rates (2018), arXiv:1708.07120 [cs.LG].
 - [31] J. Tian, N. Mithun, Z. Seymour, H.-P. Chiu, and Z. Kira, Striking the right balance: Recall loss for semantic segmentation (2022), arXiv:2106.14917 [cs.CV].
 - [32] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, Focal loss for dense object detection, in *Proceedings of the IEEE international conference on computer vision (2017)* pp. 2980–2988.
 - [33] A. Kendall, Y. Gal, and R. Cipolla, Multi-task learning using uncertainty to weigh losses for scene geometry and semantics (2018), arXiv:1705.07115 [cs.CV].
 - [34] E. L. Snider and G. Petrillo, LArSoft: Toolkit for Simulation, Reconstruction and Analysis of Liquid Argon TPC Neutrino Detectors, J. Phys. Conf. Ser. **898**, 042057 (2017).
 - [35] T. Cai, K. Herner, T. Yang, M. Wang, M. A. Flechas, P. Harris, B. Holzman, K. Pedro, and N. Tran, Accelerating Machine Learning Inference with GPUs in ProtoDUNE Data Processing, Comput. Softw. Big Sci. **7**, 11 (2023), arXiv:2301.04633 [hep-ex].
 - [36] P. Abratenko *et al.* (MicroBooNE), MicroBooNE BNB Inclusive Overlay Sample (No Wire Info), 10.5281/zenodo.8370883 (2023).
 - [37] P. Abratenko *et al.* (MicroBooNE), MicroBooNE BNB Electron Neutrino Overlay Sample (No Wire Info), 10.5281/zenodo.7261921 (2022).