Graph neural networks for position reconstruction in EXO-200

Hannah Peltz Smalley
NNPSS 2026

Background: EXO-200

- Active 2011-2018 at WIPP underground site in New Mexico
  - Phase I ended in 2014 → Phase II began in 2016
  - First observation of $2\nu\beta\beta$ decay of ^{136}Xe in 2011
  - Lower bound of $0\nu\beta\beta$ decay of ^{136}Xe: > 3.5×10^{25} y
- 200 kg LXe (≈ 130 kg of LXe in active volume)
- TPC split into two drift regions sharing a common wire grid cathode
- **Combination of scintillation and ionization signal allows full 3D reconstruction**
- **Prompt scintillation measured on two planes of Large Area Avalanche Photodiode (LAAPD)**
- Delayed ionization signal measured by crossed wires for x-y plane reconstruction

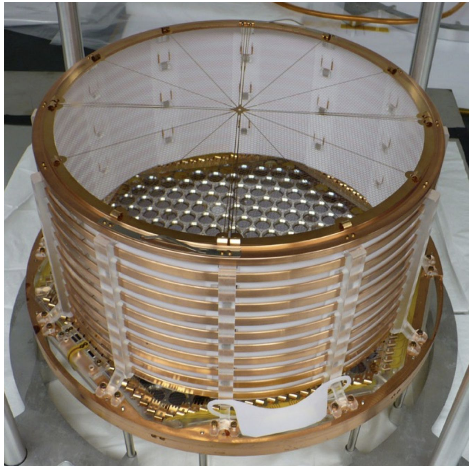

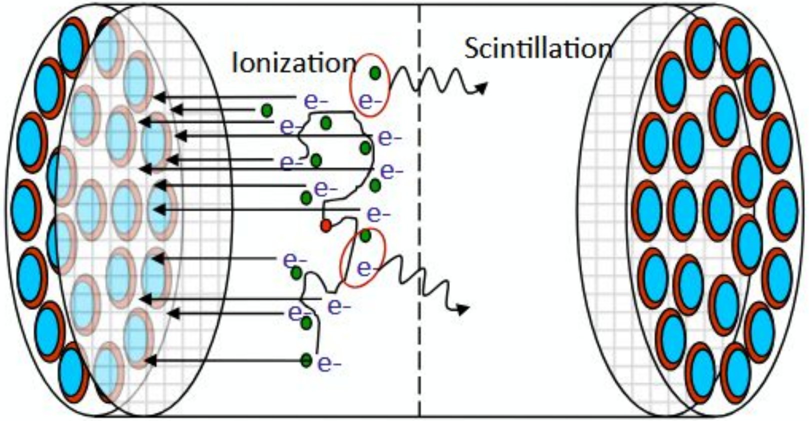

EXO-200

★ **Single-phase LXe time projection chamber**

- ○ **Prompt** **scintillation**
- ○ **Delayed, distributed** **ionization**

★ Topological discrimination between single-site (SS) and multi-site (MS) events

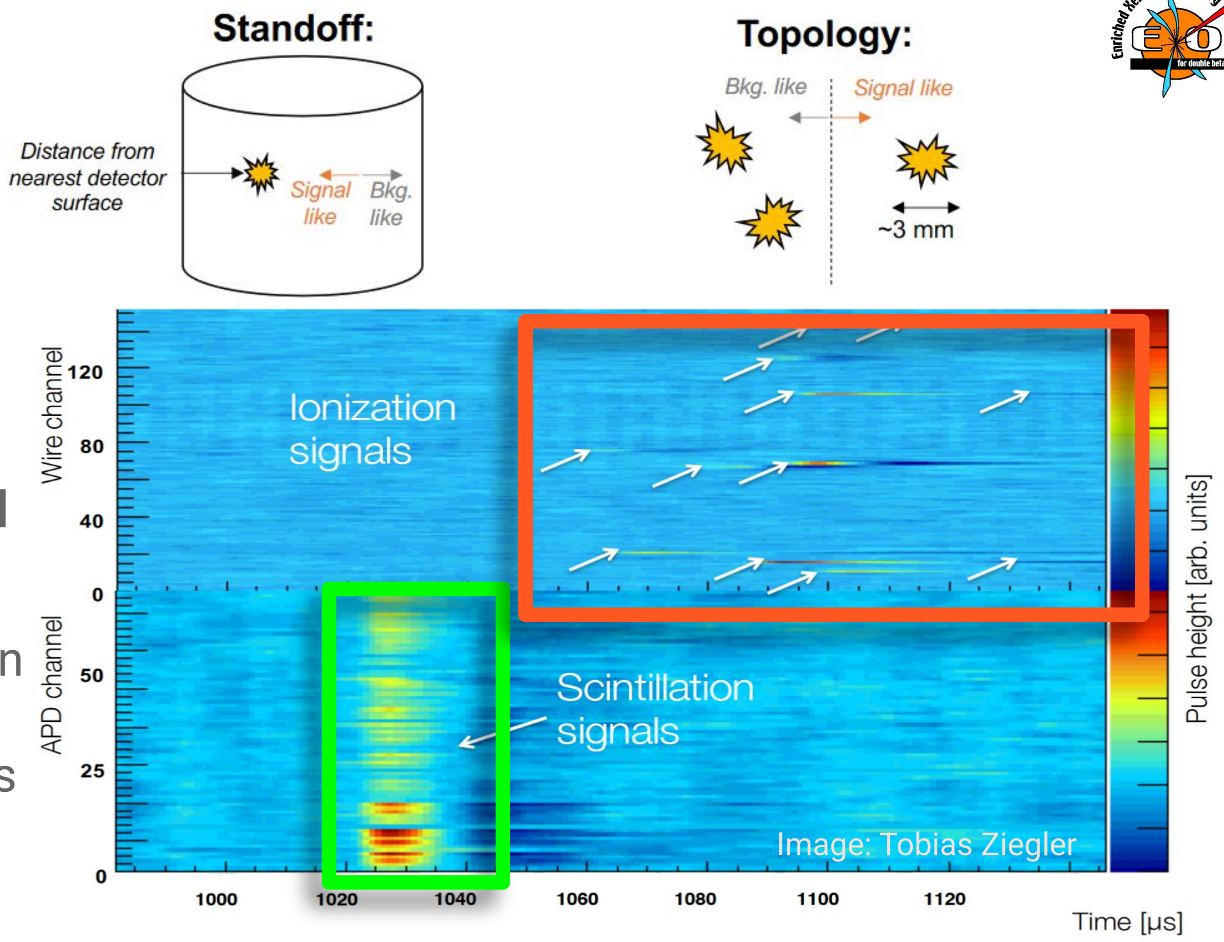

EXO-200

Traditionally, the scintillation channel is not used for position reconstruction, except in the z-coordinate.

x-y position is determined using charge collected on the crossed wire planes...

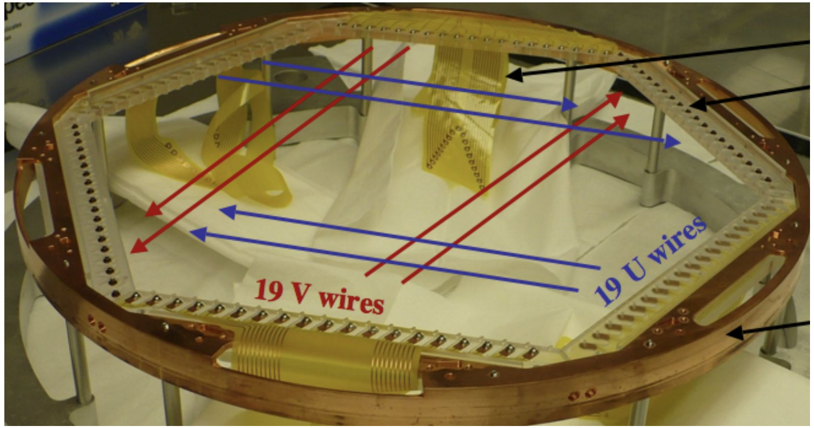

...But in events with low charge yield, one or both wire planes may not collect charge.

Alpha decays deposit most of their energy in the light channel. **Charge is insufficient to reconstruct alpha x-y position ~50% of the time.**

If we recover poorly reconstructed events, we increase available statistics for alpha-related analyses by ~2-4x. **More generally, we want to explore how the light channel can be used for position reconstruction in noble liquid TPCs.**

Previous EXO-200 light-based ML

Traditional CNNs require data to be projected onto a regular grid. Delaquis et al. stacked APD waveforms to form a heat map. The CNN applies square kernels to the image:

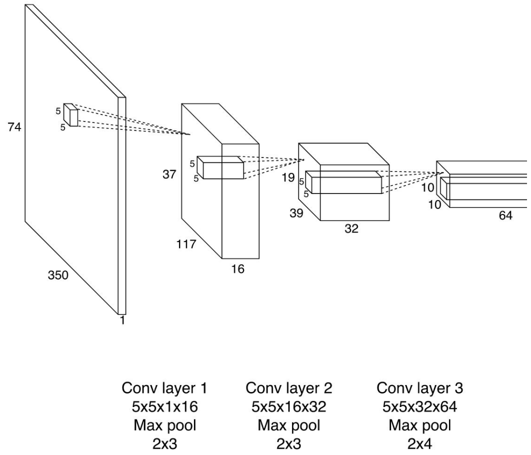

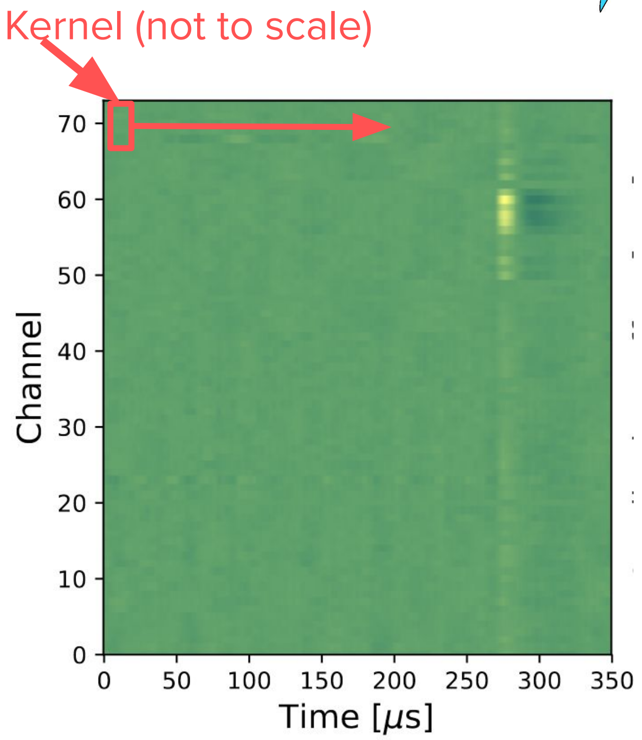

S. Delaquis et al. Deep Neural Networks for Energy and Position Reconstruction in EXO-200. JINST, 13(08):P08023, 2018.

However, EXO-200 APD planes are hexagonal.

Making this geometry CNN-compatible requires the geometry to be projected as a 2D rectangle, which can be lossy.

But projecting the hexagonal hit map would introduce empty pixels, which are computationally wasteful.

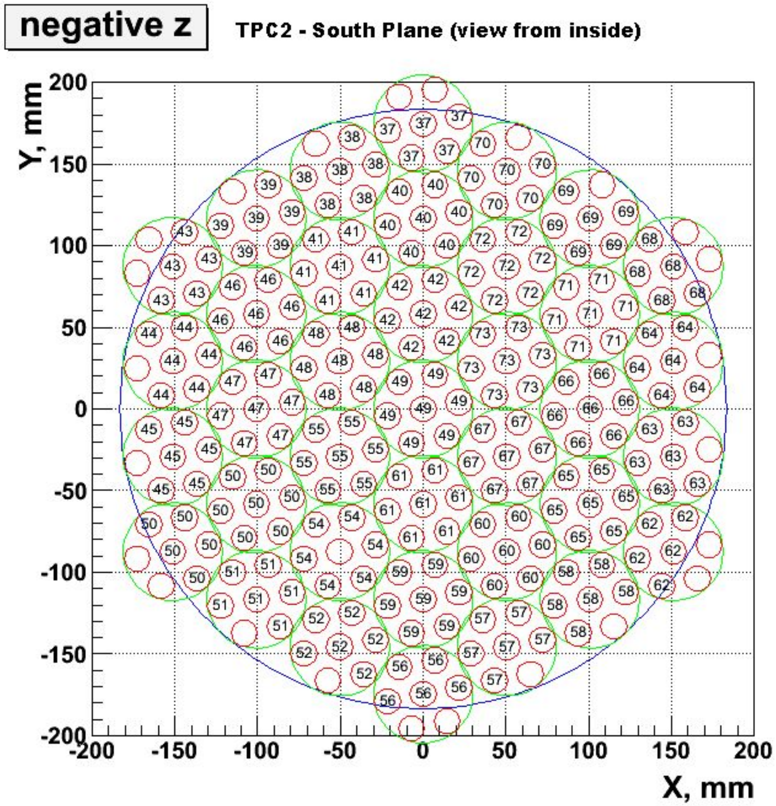

New approach: Graph Neural Network

Represent APD channels as nodes in a graph, connected by edges.

Convolution operations are generalized to apply over connected nodes.

Graph neural networks are relatively new; graph convolutional networks were formulated by Kipf & Welling in 2017.

They are especially useful when handling complicated detector geometries and heterogeneous data types.

Graph construction

Each APD channel is a node with one feature: **the number of photon counts**

Connected to its nearest neighbor nodes + its opposite plane node by edges

Edge weights scale the strength of the message from one node to another

Fixed edge weights based on $1/r^2$ distance between APD channels

Network does not know the position of each node; only weights & adjacent connections!

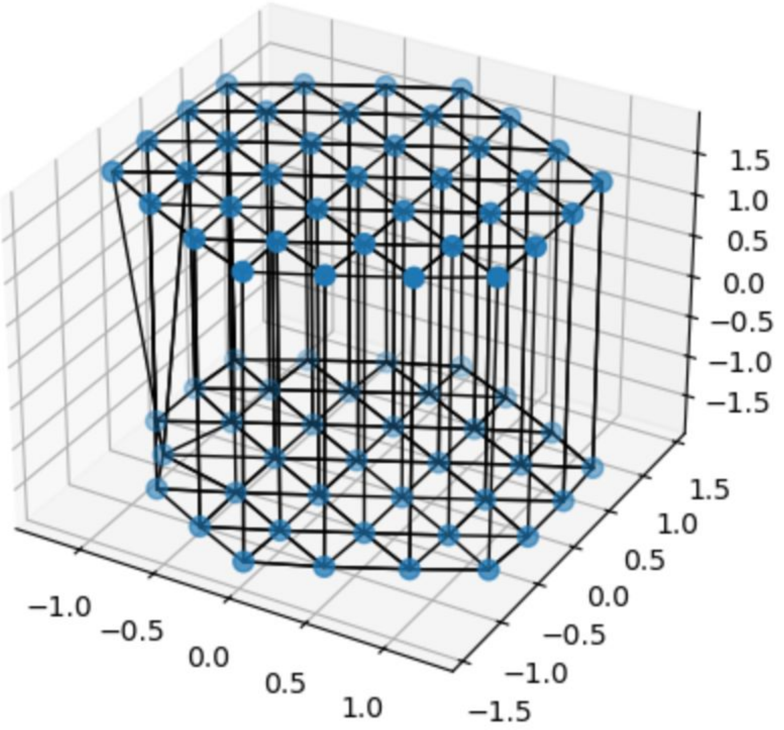

Example Neighbor Aggregation Operation (dim 1)

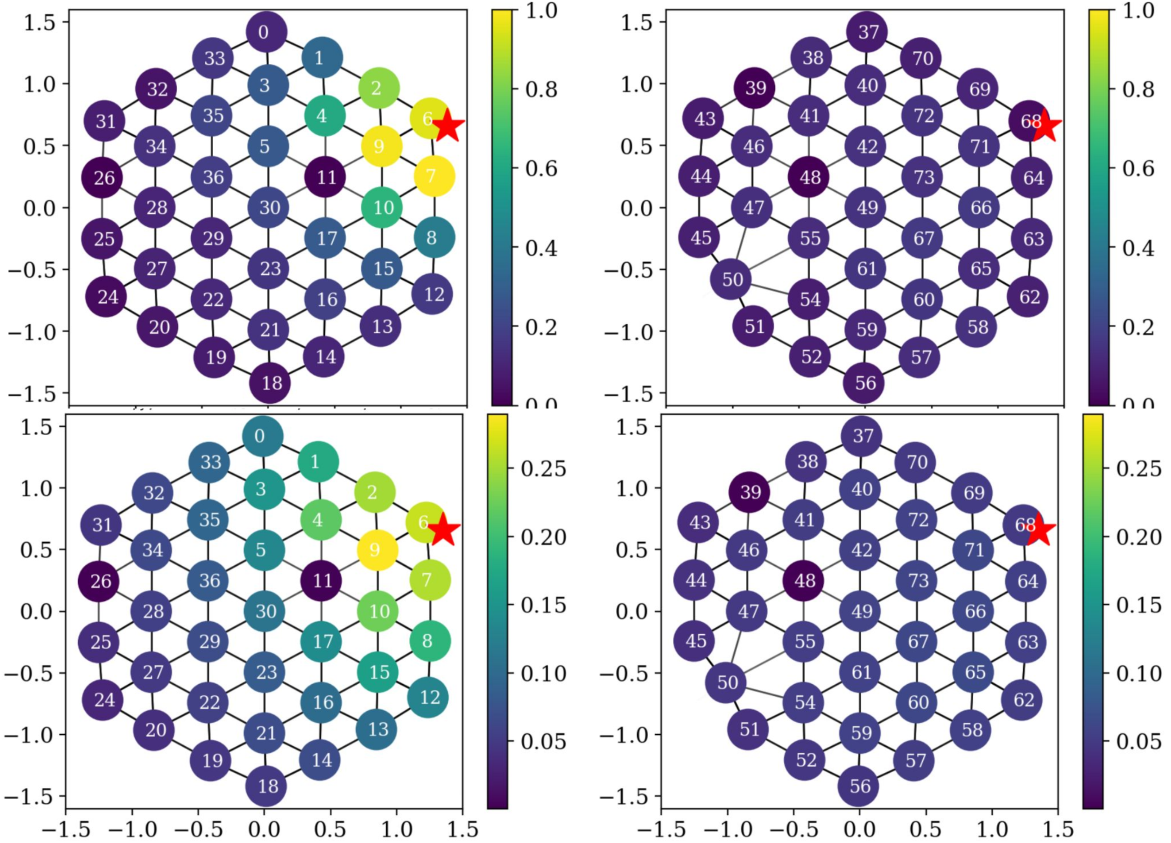

Result of 4 convolution operations

“Brightest” nodes are enhanced

Training Set

Model was trained on alpha decays from an end-of-run calibration campaign, where ^{222}Rn was injected into the detector volume.

EXO-200 does not simulate APD hits, so the model was trained on reconstructed positions from the charge channel as “truth” values.

Thus, only events for which full charge reconstruction was possible could be used in training.

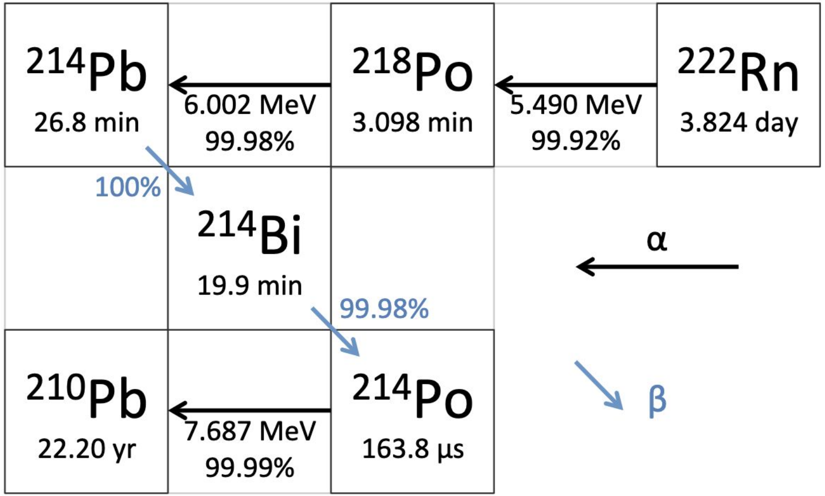

Model Structure

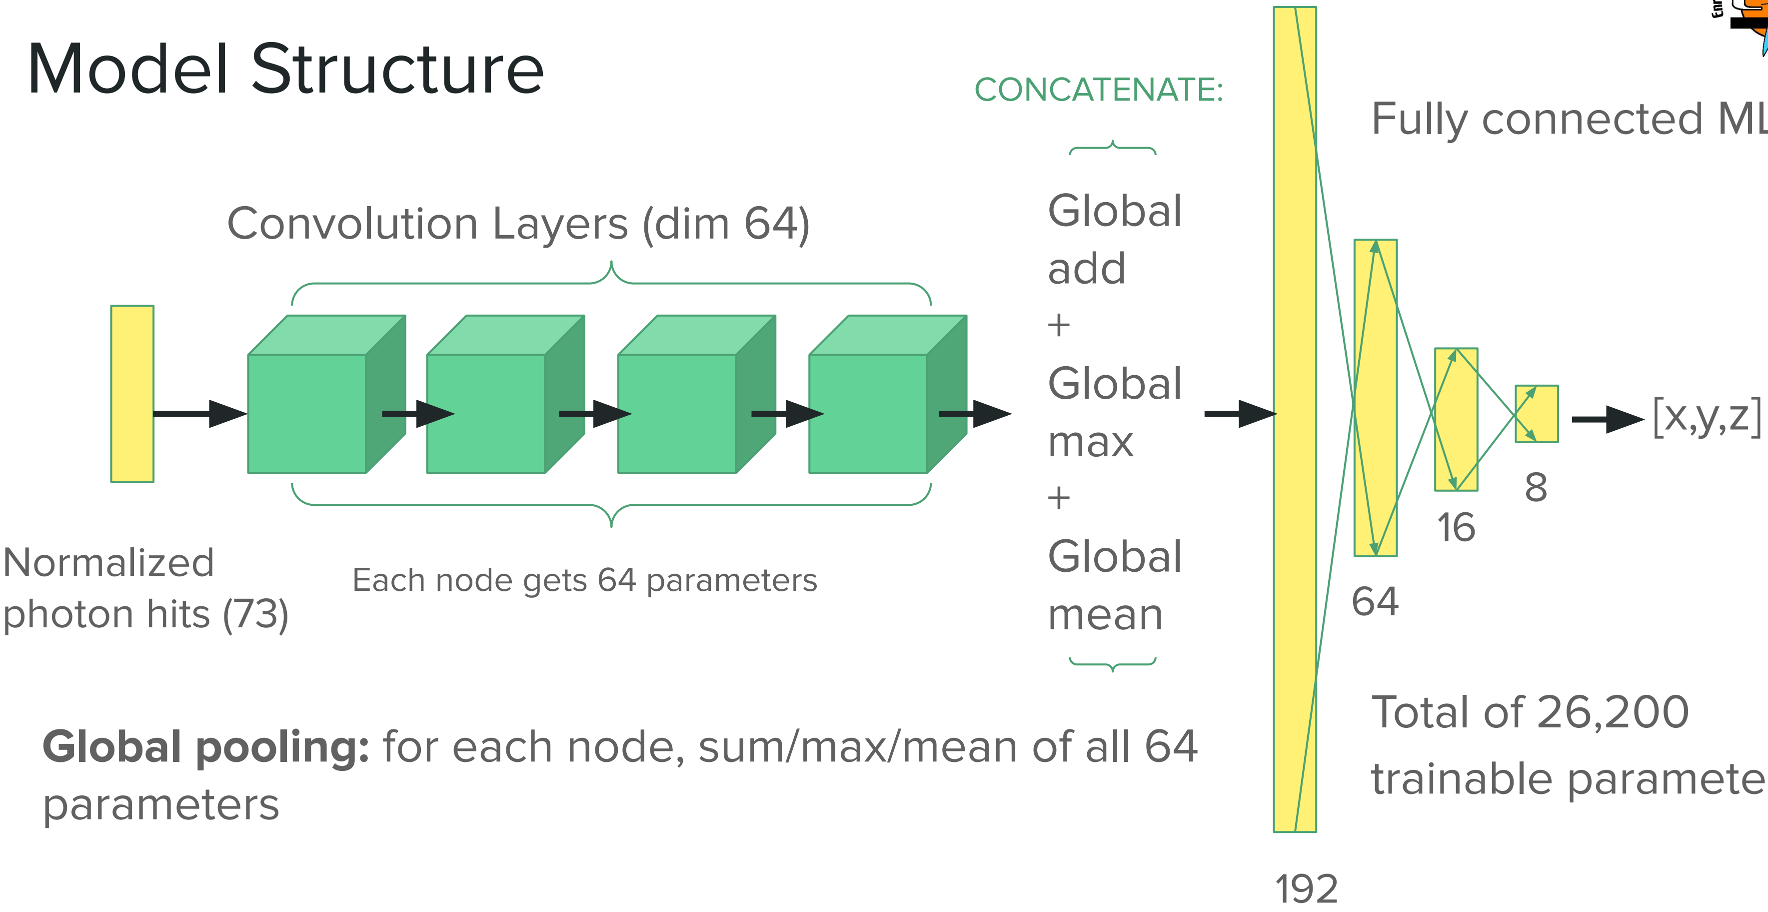

Global pooling: for each node, sum/max/mean of all 64 parameters

Model Training

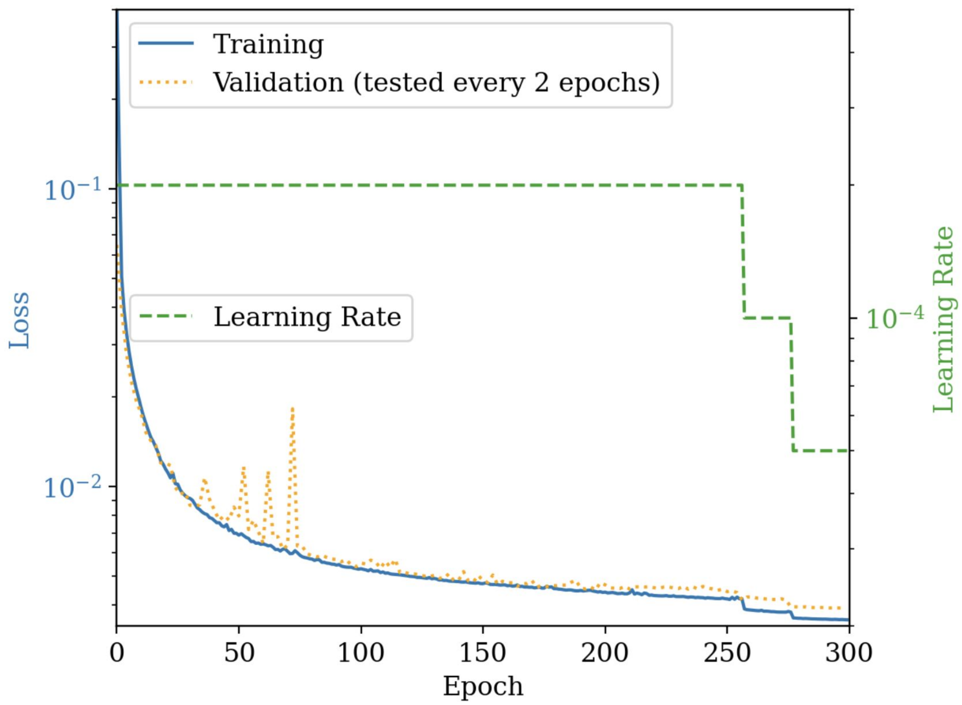

Model was trained for 300 epochs with a variable learning rate at batch size of 32.

The learning rate was initially 2e-4 and automatically reduced by half when detecting a plateau.

Total training time was about 12.3 hours for 300 epochs.

Preliminary Results

Accuracy: ^{222}Rn Chain

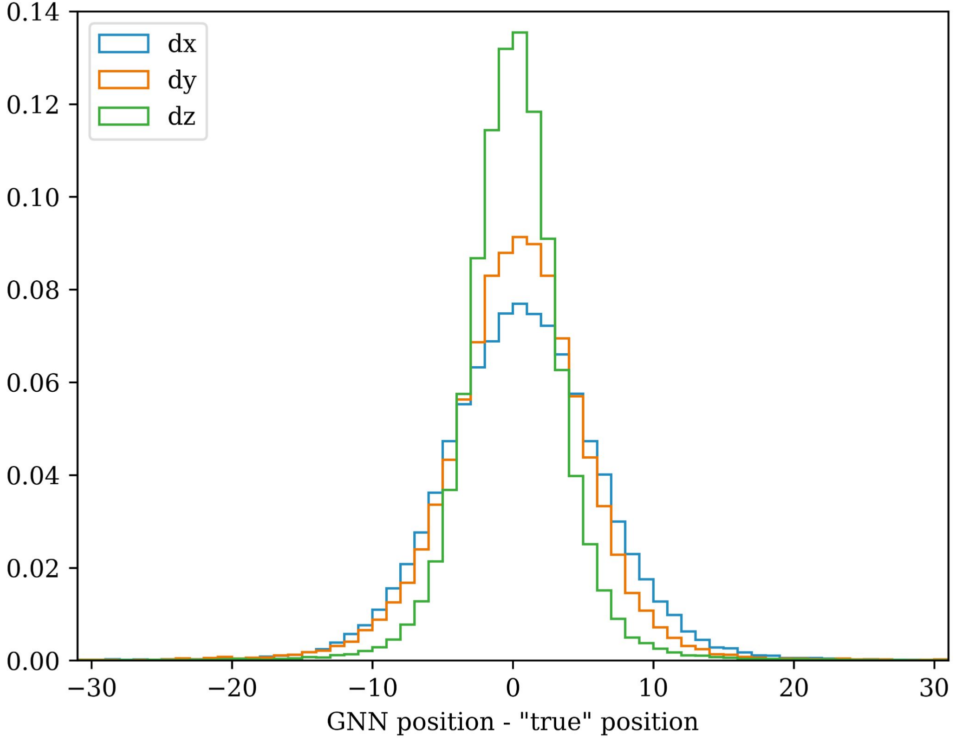

- mean squared residual in x: 4.92 mm
  - standard deviation: 9.03 mm
- mean squared residual in y: 4.04 mm
  - standard deviation: 6.5 mm
- mean squared residual in z: 2.76 mm
  - standard deviation: 5.65 mm

These are all improved over the CNN!

Accuracy: ^{222}Rn Chain

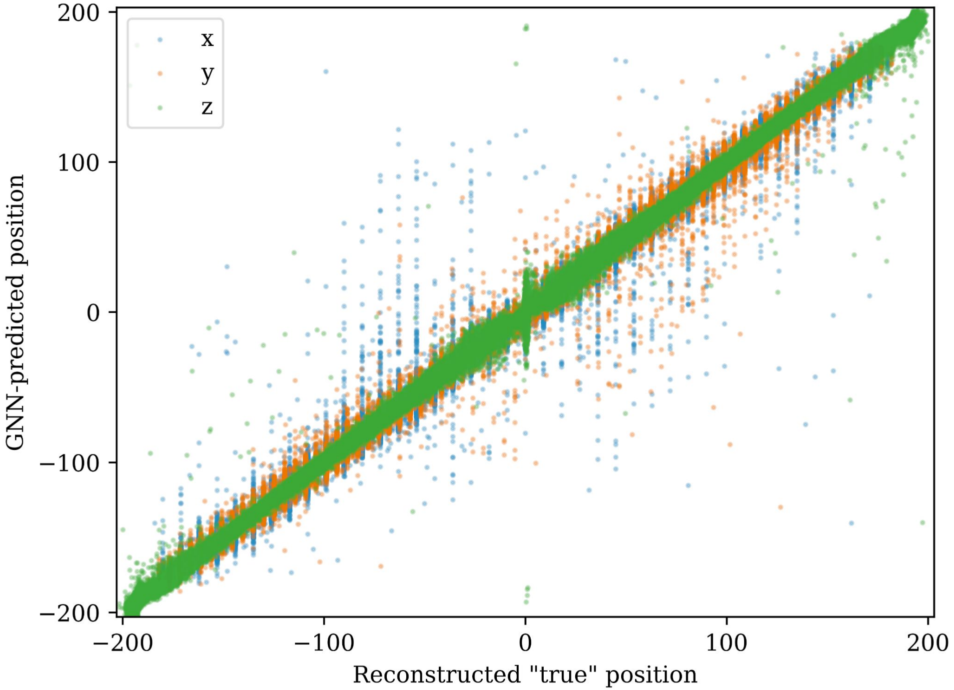

The GNN is well-correlated with the traditional reconstruction. Most outliers are in the x-direction, possibly due to the wire pitch

The most problematic area is the cathode. Likely some ambiguity in the traditional reconstruction + detector effects in this region

Accuracy: ^{222}Rn Chain

The GNN reproduces the distribution of events in x, y, and z.

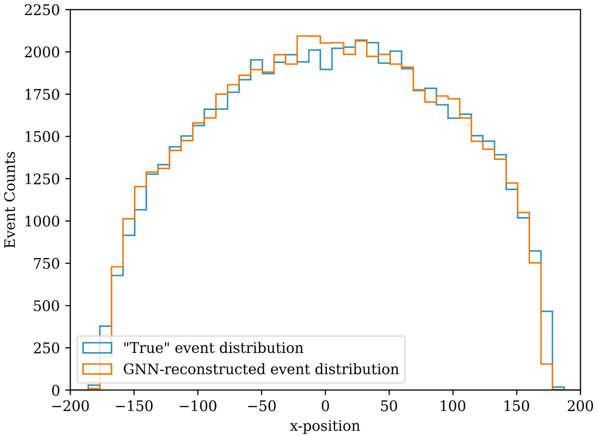

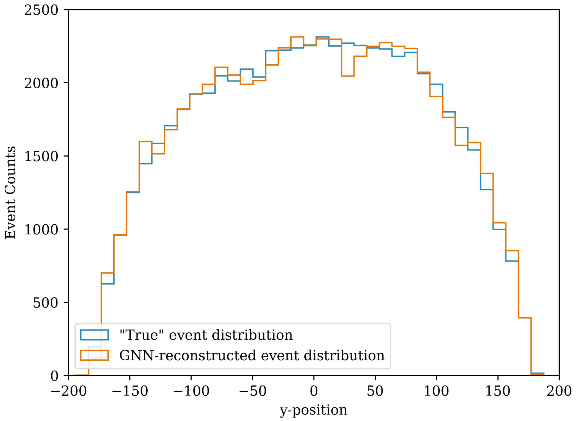

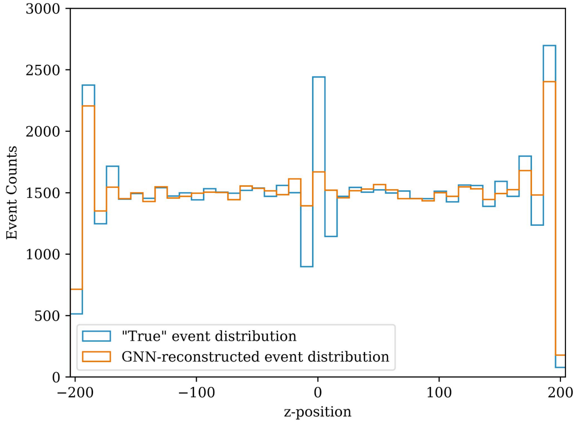

Conclusions

- The model **performs well on alpha events** in the ^{222}Rn chain!
  - Accurate to order ~10 mm
  - Also shown to work on ^{220}Rn chain events
- However, other tests show **it doesn't work very well on low-scintillation events** like betas and gammas (it was not trained on these)
- Since a full light simulation is not available, the model can't be trained or tested on simulated events
  - Future detectors will have more advanced light detection systems and simulations
- We are considering whether to apply the model in situations where highly exact position reconstruction is not required, e.g. matching ^{222}Rn decays to progeny ^{218}Po decays.

Backup Slides

EXO-200

Number of electrons and photons from an event is anti-correlated and depends on electric field

★ Larger E-field → more ionization

β, γ events deposit light + charge quanta characterized by θ^R

$$E_R = E_S \cdot \sin(\theta^R) + E_I \cdot \cos(\theta^R)$$

“Rotated energy” = linear combination of light and charge

θ^R measured every week with ^{228}Th source

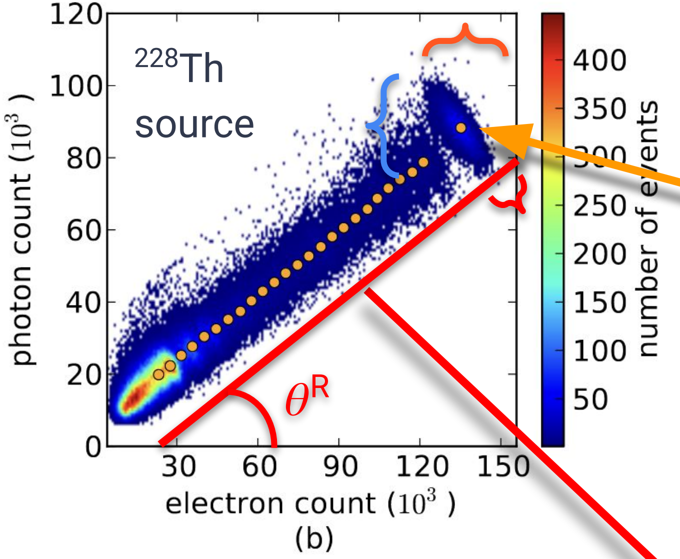

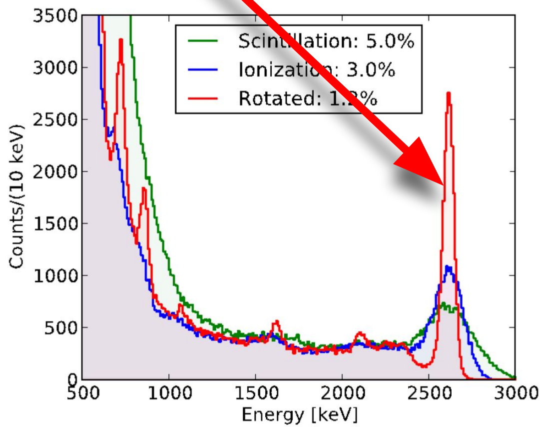

Highly simplified example

Consider a plane of 5 detectors, each represented by a vector: $$\mathbf{x}_i = [x_i, y_i, z_i, A_i]$$

Edges connect nearest neighbors: D0 is connected to all detectors. Outer detectors connect only to D0.

GNN first maps raw features to an embedding space via a weight matrix. Dimension of embedding space is a hyperparameter. For readability, assume embedding dimension = 3.

At some training step, weight matrix looks like:

	D3 (0,1,0,9.0)	
D1 (1,0,0,8.0)	D0 (0,0,0,12.0)	D2 (-1,0,0,6.0)
	D4 (0,-1,0,5.0)	

*completely arbitrary numbers!

$$W = \begin{bmatrix} 0.6 & 0.0 & 0.0 & 0.1 \\ 0.0 & 0.6 & 0.0 & 0.1 \\ 0.0 & 0.0 & 0.2 & 0.4 \end{bmatrix}$$

Highly simplified example

Mapping the features of D0 to the embedding:

$$\mathbf{h}_{D0}^{(0)} = W\mathbf{x}_{D0} = \begin{bmatrix} 0.6 * 0 + 0.1 * 12 \\ 0.6 * 0 + 0.1 * 12 \\ 0.2 * 0 + 0.4 * 12 \end{bmatrix} = \begin{bmatrix} 1.2 \\ 1.2 \\ 4.8 \end{bmatrix}$$

To incorporate information from the surrounding detectors, we have to do the same on D1, D2, D3 and D4, then aggregate the result via element-wise mean:

Detector	$\mathbf{h}^{(0)}$
D1	[1.4, 0.8, 3.2]
D2	[0.0, 0.6, 2.4]
D3	[0.9, 1.5, 3.6]
D4	[0.5, -0.1, 2.0]

$$\mathbf{m}_{D0} = \begin{bmatrix} 0.7 \\ 0.7 \\ 2.55 \end{bmatrix}$$

Highly simplified example

Update the embedding of D0 with its neighbors' features:

$$\mathbf{h}_{D0}^{(1)} = \mathrm{ReLU}(\mathbf{h}_{D0}^{(0)} + \mathbf{m}_{D0}) = \mathrm{ReLU}(\begin{bmatrix} 1.9 \\ 1.9 \\ 7.35 \end{bmatrix}) = \begin{bmatrix} 1.9 \\ 1.9 \\ 7.35 \end{bmatrix}$$

This is the embedding of D0 after one GNN layer.

After some number of layers, embeddings can be input to an MLP to predict position.

At the end of each training loop, the prediction is compared to the true event position and a loss function is computed. The gradients of the loss are computed with respect to all weights. An optimizer updates the weight matrix in response.

Weight Initialization

Default PyTorch initialization (Kaiming uniform)

e.g. normalized hits on node ➜ {64 uniformly distributed random numbers}

The weights are then updated by the Adam optimizer during training.

There are a total of 26,200 trainable parameters (compare to 2.4e6 in 2018 DNN).

Summary: Model Structure

- **Input:** each node initially has 1 feature (photon hits on channel)
  - Initialize 64 hidden parameters per node.
- **Convolution layers:** perform one nearest-neighbor sum aggregation operation per layer
  - literature shows GNNs perform best with 4 or fewer layers.
- **Global pooling:** 3 pooling operations from the 64 parameters.
  - Using add, max, and mean and concatenating the results gives the network a balanced picture of the features.
- **MLP:** Reduce 192-length tensor to 3 coordinates using fully-connected layers.

Accuracy

"Benchmark" using the results from the 2018 DNN paper:

Their results were very dependent on the spatial distribution of the test data. Their "best case" scenario is a tight fiducial cut

In this scenario the DNN achieved the following mean residuals:

dx = 8.3 mm, dy = 4.8 mm, dz = 6.2 mm

Goal is for the GNN to meet or exceed this.

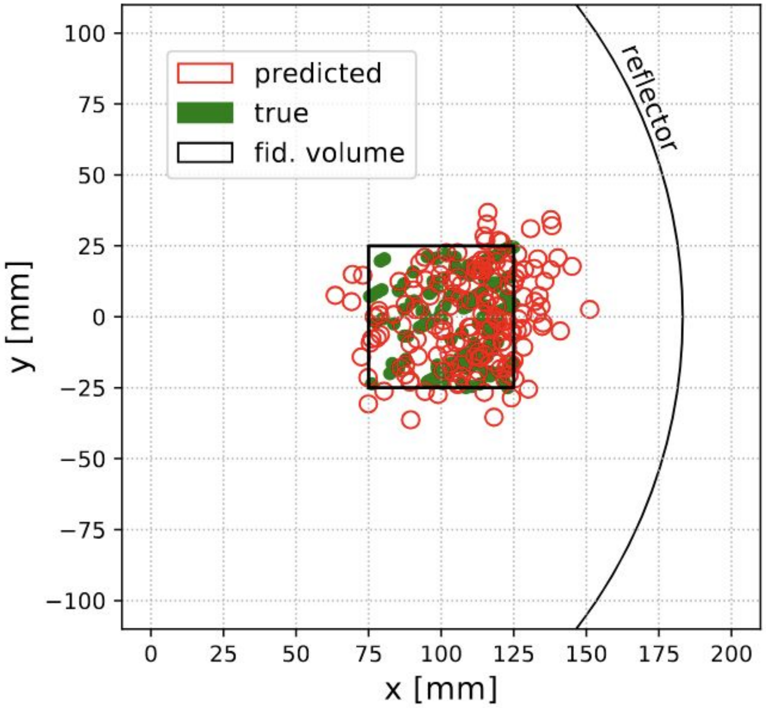

Accuracy: low-background alphas

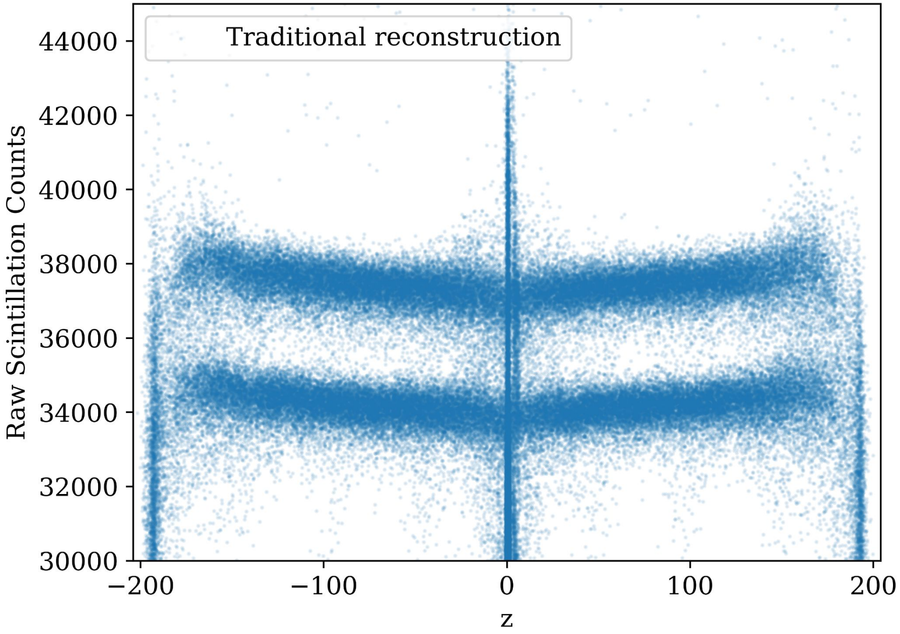

Looking at the scintillation energy plotted against z-position, there is an interesting effect in the traditional reconstruction which appears to skew towards the anodes.

These are likely surface alphas.

The GNN reconstructs these events more diffusely. There may be an optical or electric field effect in this region

Accuracy: low-background alphas

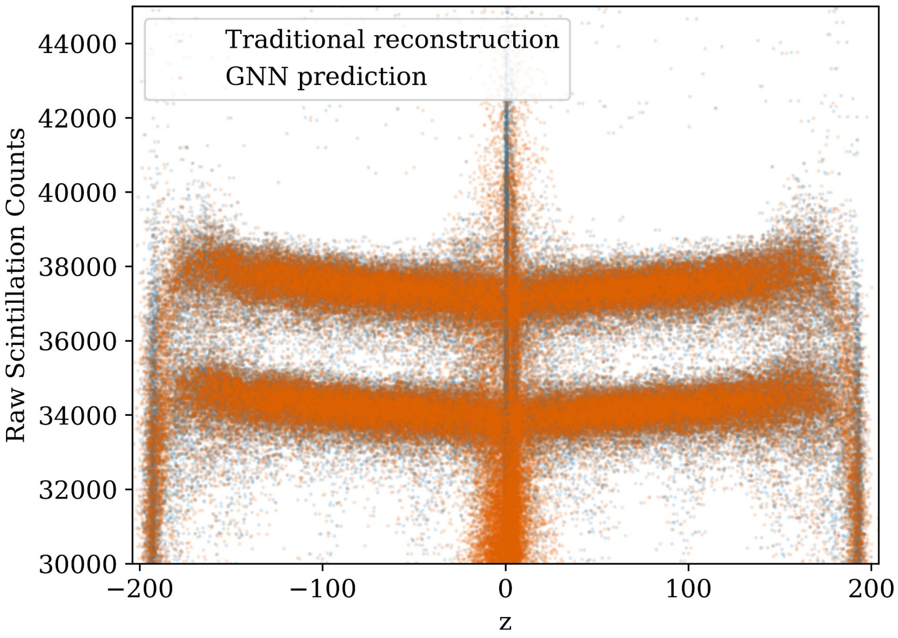

Looking at the scintillation energy plotted against z-position, there is an interesting effect in the traditional reconstruction which appears to skew towards the anodes.

These are likely surface alphas.

The GNN reconstructs these events more diffusely. There may be an optical or electric field effect in this region

Accuracy: low-background alphas

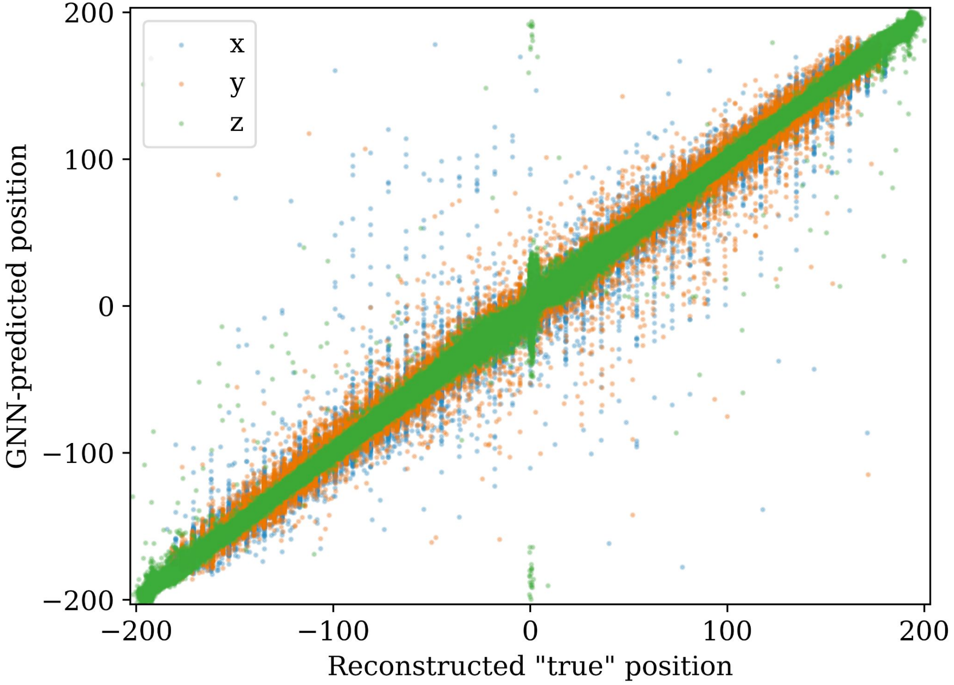

The model performed almost identically on the LB data as the ^{222}Rn campaign data. This is good news

Accuracy: Th source calibration data (cathode position)

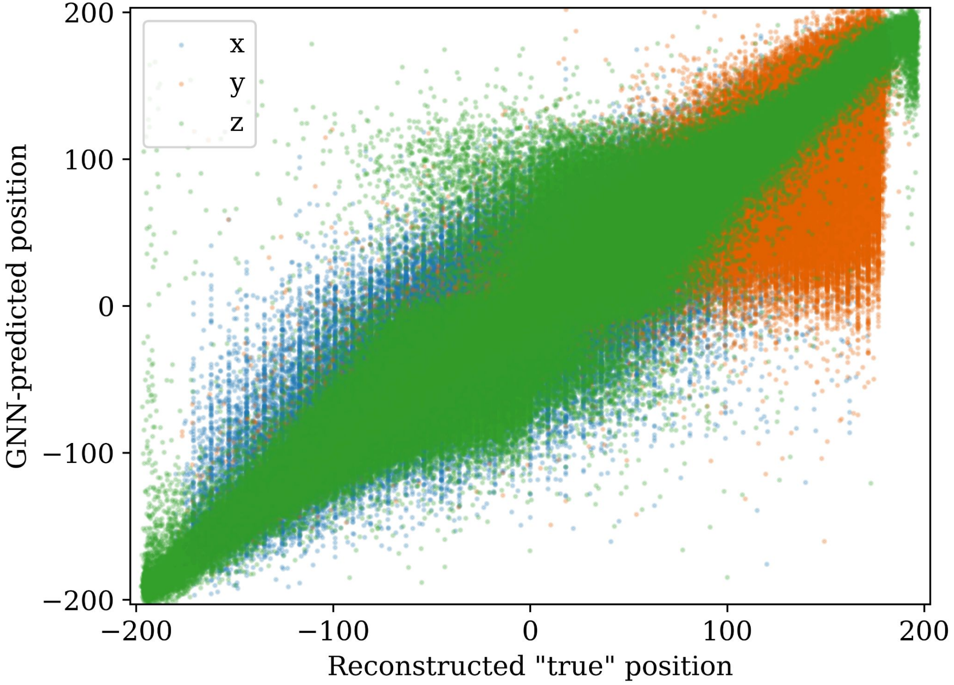

Accuracy on this dataset is not good and is only marginally improved by a fiducial z cut.

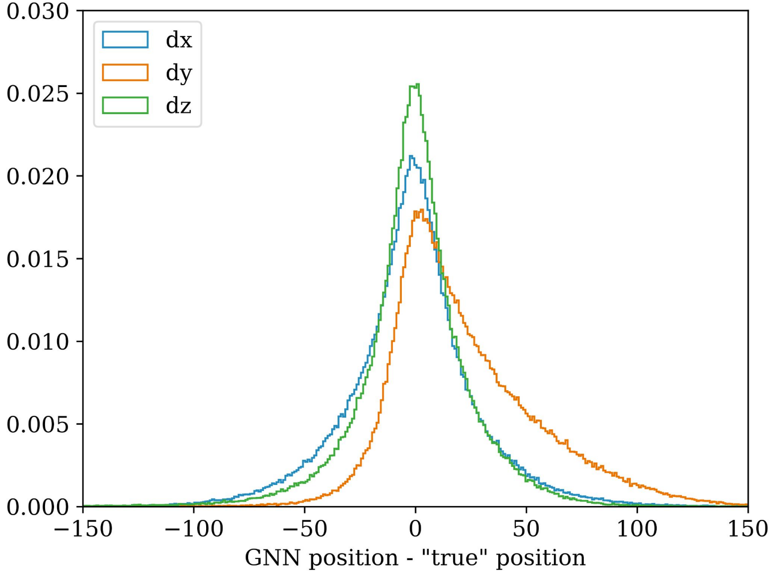

Accuracy: Th source calibration data (cathode position)

The discrepancy in the y-direction was also seen in the 2018 paper. This may be unrelated to how well the model is trained and reflect ambiguities in the traditional reconstruction near detector edges

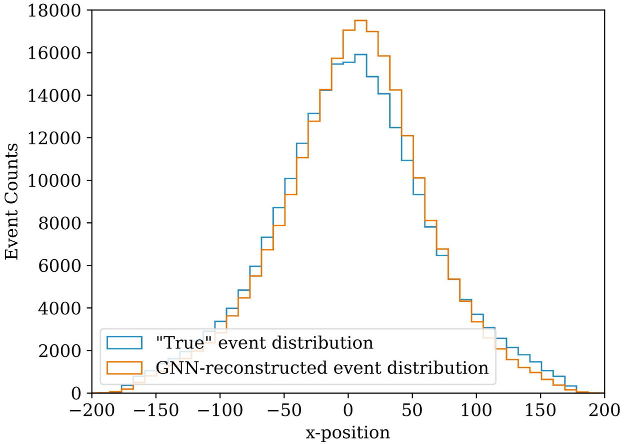

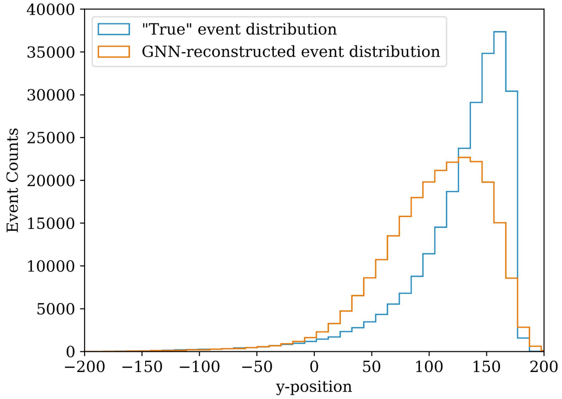

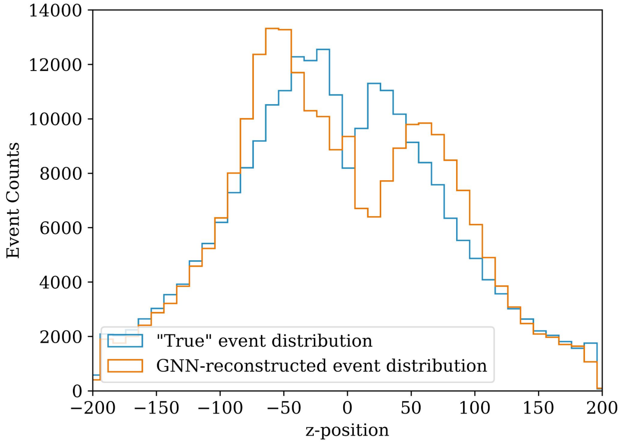

Accuracy: ^{220}Rn Campaign

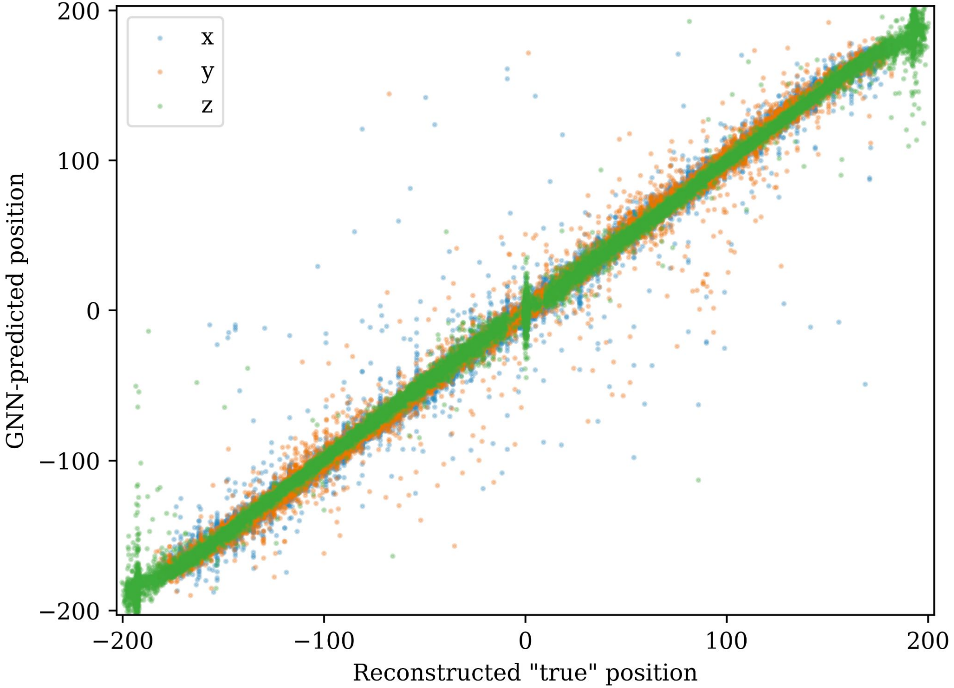

This plot includes alpha and alpha-beta (BiPo) events as well as uncategorized alpha events in the dataset.

The GNN is generally well-correlated with the traditional reconstruction, but has a skew towards the anodes

Most of the skew comes from including alphas with energy somewhere between the single Bi alpha and the BiPo

These events occur close to the anodes and appear to be alphas that have lost some energy

Accuracy: ^{220}Rn Campaign

The GNN tends to place events farther from the anodes than the traditional reconstruction.

Events were isotropically sampled in the z-direction to get a more uniform distribution.

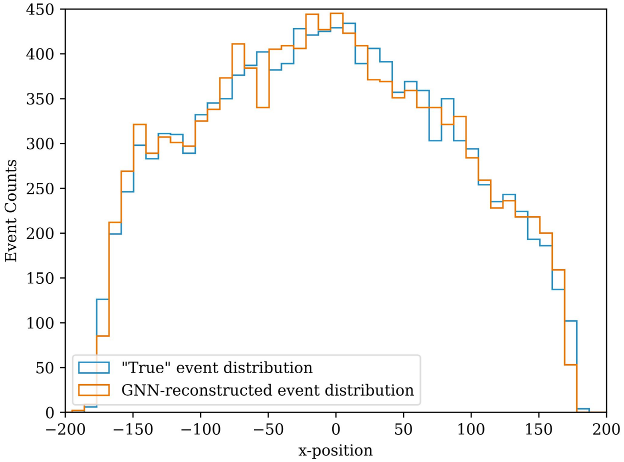

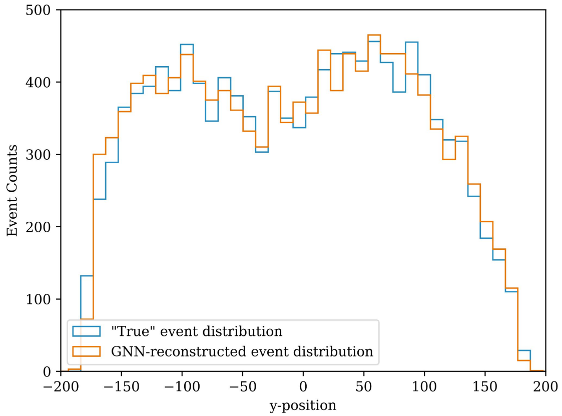

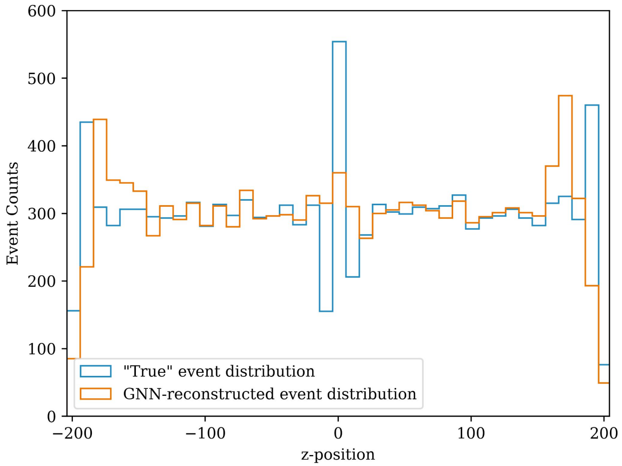

Accuracy: ^{220}Rn Campaign

Without any cuts except scintillation counts > 30,000:

- mean squared residual in x: 5.08 mm
  - standard deviation: 9.47 mm
- mean squared residual in y: 4.5 mm
  - standard deviation: 7.3 mm
- mean squared residual in z: 3.79 mm
  - standard deviation: 6.46 mm

Despite the issues with z, the result is strong.