AI -driven Reconstruction Pipelines in Neutrino Experiments



Jianming Bian

University of California, Irvine

10-28-2026

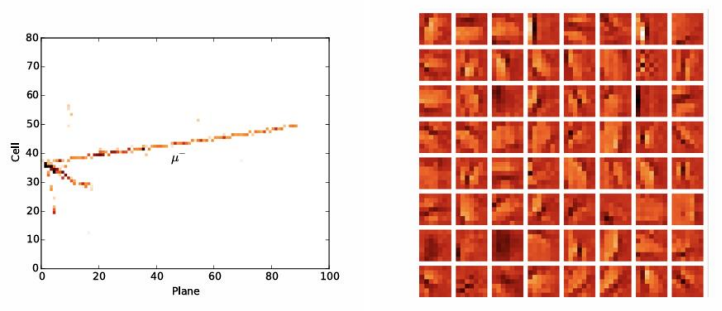
NNN2026, Chung-Ang University, Seoul

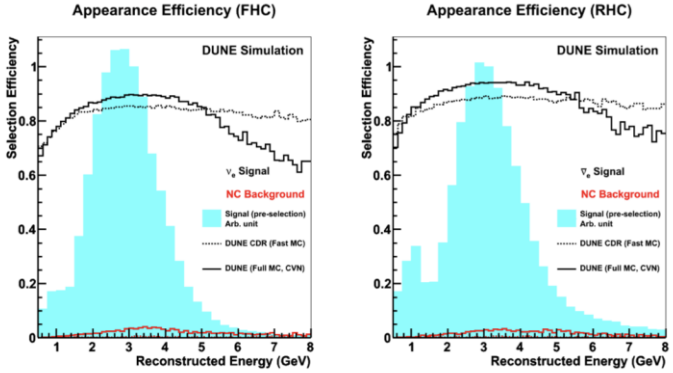
Introduction

- Reconstruction Tasks
  - Signal Processing and Hit Finding
  - Clustering
  - Kinematics
  - Event-Classification and Particle Identification
- AI architectures for reconstruction
  - CNN / Sparse CNN: Local feature extraction, ROI detection, segmentation, PID, kinematic regression
  - GNN: Sparse and irregular hits, particle clustering, association, interaction reconstruction
  - Transformer: attention based global correlations, multi-view fusion, clustering, joint reconstruction
- Emerging Directions
  - Foundation Models
  - Multimodal AI
  - Agentic AI

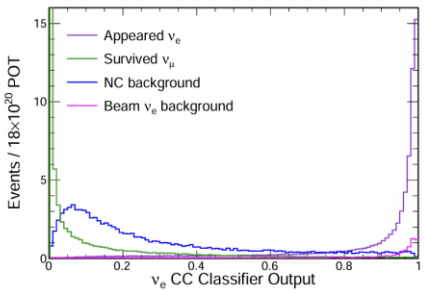
CNN for PID and Event Classification

- CNNs extract local spatial features and learn hierarchical representations through convolutional filters.
- CNNs are widely used in neutrino experiments for event classification, particle identification (PID), and kinematic reconstruction.

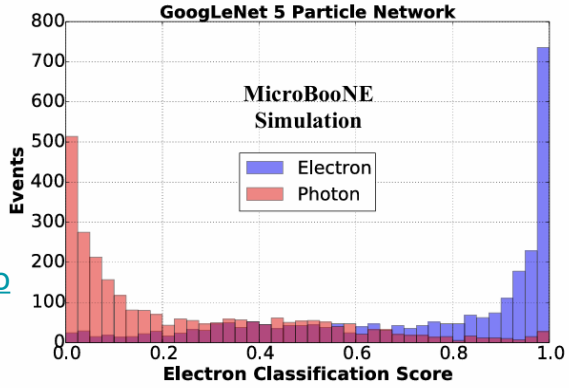




Neutrino Interaction Classification with a Convolutional Neural Network in the DUNE Far Detector, Phys. Rev. D 102, 092003 (2020, DUNE)



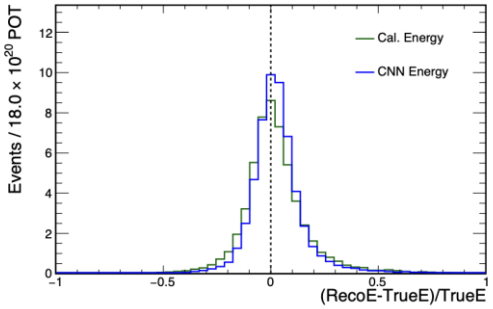
Aurisano et al., A Convolutional Neural Network Neutrino Event Classifier, JINST 11, P09001 (2016, NOvA)



Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber, JINST 12, P03011 (2017, MicroBooNE)

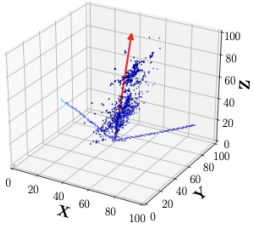
CNN for Kinematics Reconstruction

Electron Shower Energy

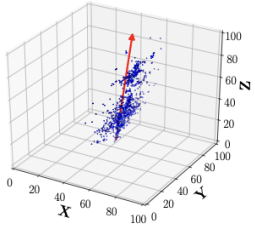


Baldi et al., Improved Energy Reconstruction in NOvA with Regression CNNs, Phys. Rev. D 99, 012011 (2019)

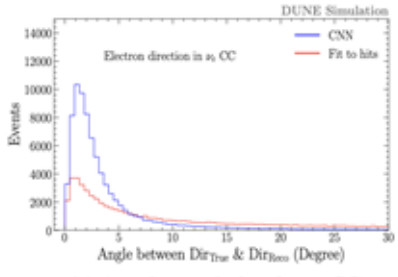
Track/Shower Direction



(a) Full-event ν_e CC



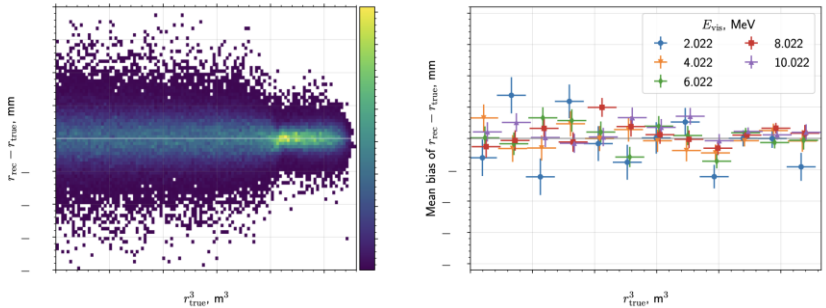
(b) Prong-only ν_e CC



(a) Angular resolution for ν_e CC

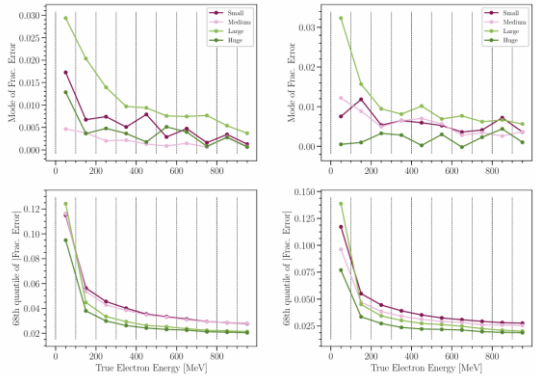
Liu et al., Deep-Learning-Based Kinematic Reconstruction for DUNE, NeurIPS Workshop (2020)

Vertex



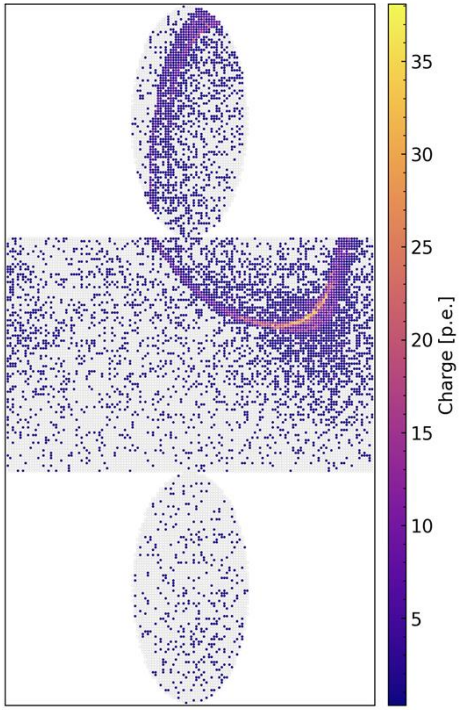
Qian et al., Vertex and Energy Reconstruction in JUNO with Machine Learning Methods, NIM A 1010, 165527 (2021)

Electron Shower Energy

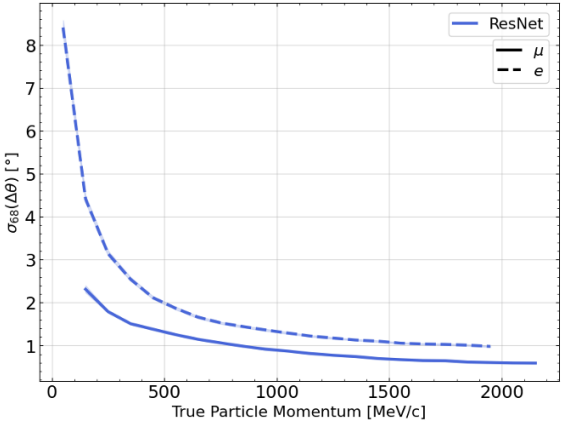


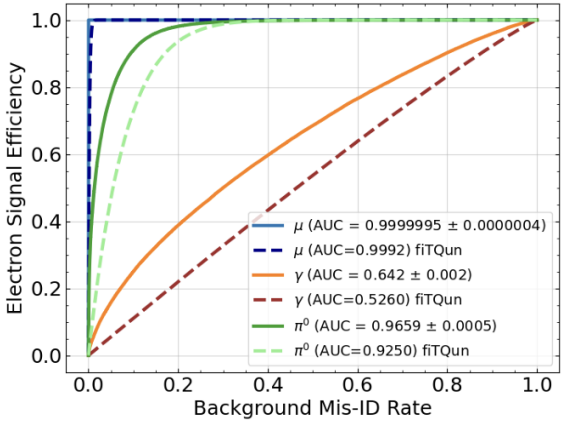
Carloni et al., Convolutional Neural Networks for Shower Energy Prediction in Liquid Argon Time Projection Chambers, JINST 17, P02022 (2022)

ResNet-based Reconstruction: PID + Kinematics



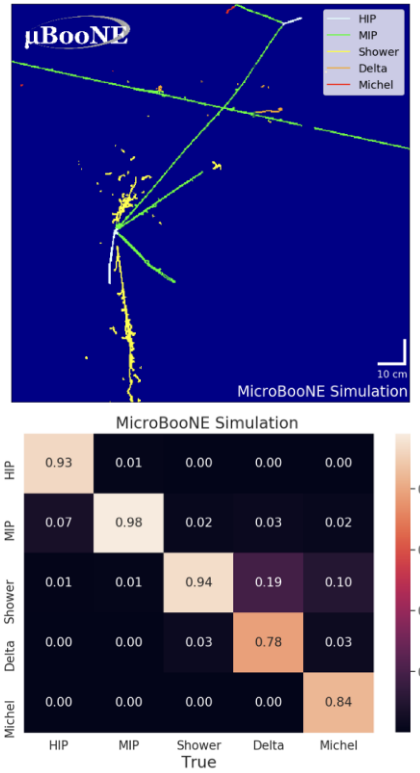
Vertex, direction, momentum, PID



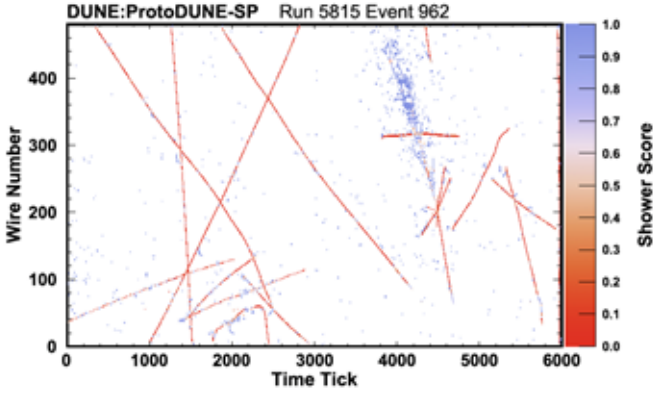


Atta et al., Enhancing Event Reconstruction in Hyper-Kamiokande with Machine Learning: A ResNet Implementation, arXiv:2604.13503 (2026)

CNN for Semantic Segmentation



Abratenko et al. (MicroBooNE Collaboration), Semantic Segmentation with a Sparse Convolutional Neural Network for Event Reconstruction in MicroBooNE, Phys. Rev. D 103, 052012 (2021, MicroBooNE)



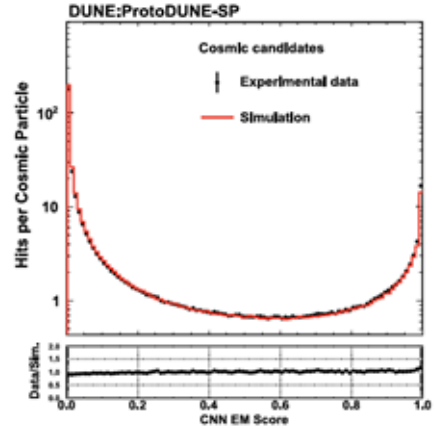


Fig. 9 The CNN shower classifier scores for cosmic-ray muon hits from experimental data (black) and simulation (red). The error bars on the data are statistical

DUNE Collaboration, Separation of Track- and Shower-like Energy Deposits in ProtoDUNE-SP Using a Convolutional Neural Network, Eur. Phys. J. C 82, 903 (2022, ProtoDUNE-SP)

Reconstruction Framework: Wire-Cell

Wire-Cell is a scalable, cell-based signal processing and 3D event reconstruction framework for LArTPC detectors, extensively used in MicroBooNE, with components deployed or developed for DUNE, ProtoDUNE, SBND, and ICARUS. It integrates machine learning for signal processing (U-Net CNN), neutrino vertex reconstruction (Sparse CNN), and electron-neutrino identification (BDT/XGBoost), while particle clustering and tracking remain primarily physics-driven.

- Scalable Inference for LArTPC Signal Processing with MobileU-Net and Overlap–Tile Chunking https://ml4physicalsciences.github.io/2025/files/NeurIPS_ML4PS_2025_280.pdf
- CNN (U-Net) for regions of interest (ROI) detection in DUNE signal processing algorithm WireCell, JINST https://iopscience.iop.org/article/10.1088/1748-0221/16/01/P01036

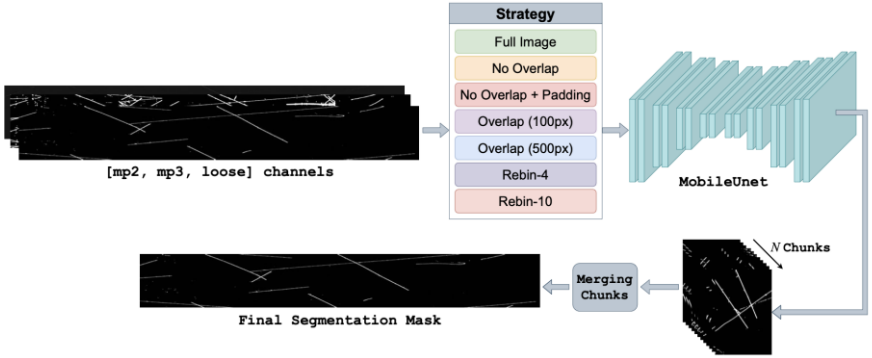


Figure 2: Schematic diagram of the proposed approach.

Table 1: Comparison of U-Net and MobileU-Net across inference strategies on the LArTPC ROI task. Metrics: Intersection over Union (IoU), Purity, Efficiency on the test set; single-instance CPU memory and inference time. The definition of each strategy can be found in Sec 2.4.

		Test Set			Single Instance	
Model	**Strategy**	**IoU**	**Purity**	**Efficiency**	**Memory (MB)**	**Time (s)**
U-Net	Full image	0.83	0.89	0.88	6193.7	106.57
	No Overlap	0.82	0.89	0.87	936.2	100.47
	No Overlap + Padding	0.83	0.83	0.88	1028.5	126.96
	Overlap (100 px)	0.83	0.89	0.87	954.4	132.71
	Overlap (500 px)	0.83	0.89	0.87	954.5	552.28
	Rebin 4	0.55	0.52	0.83	1811.8	30.33
	Rebin 4 (Retrained)	0.79	0.83	0.88	1812.8	30.64
	Rebin 10	0.31	0.34	0.48	918.9	12.17
MobileU-Net	Full image	0.84	0.89	0.88	1719.7	14.41
	No Overlap	0.77	0.82	0.87	442.4	12.90
	No Overlap + Padding	0.75	0.85	0.84	463.5	12.85
	Overlap (100 px)	0.80	0.86	0.87	442.4	12.08
	Overlap (500 px)	0.82	0.88	0.88	442.4	50.28
	Rebin 4	0.63	0.44	0.85	635.0	4.29
	Rebin 4 (Retrained)	0.83	0.86	0.88	635.0	4.32
	Rebin 10	0.48	0.33	0.61	435.5	2.39

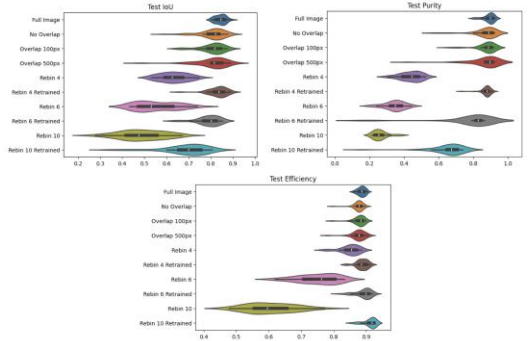


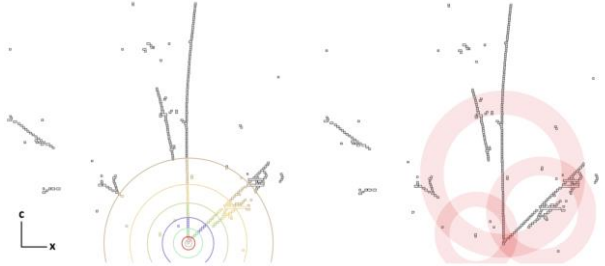
Figure 3: MobileUnet's IoU, Purity, and Efficiency distribution violin plots over the test set via multiple strategies.

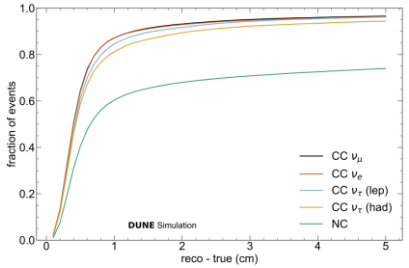
Reconstruction Framework: Pandora

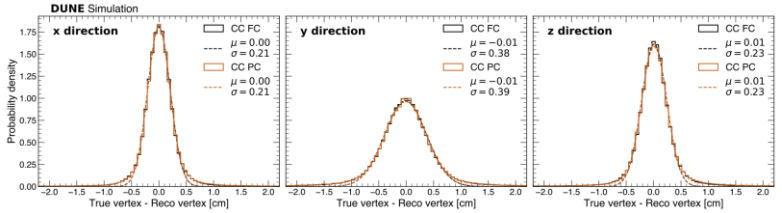
Pandora: A widely used particle reconstruction framework for LArTPC detectors, including MicroBooNE, ProtoDUNE, DUNE, SBND, and ICARUS, with increasing integration of AI/ML techniques in recent years.

CNN-based vertex reconstruction

- Predicts distance to the vertex position for each hit and uses ring-drawing method to predict vertex position
- Published last year, and a new DUNE preprint describing its use in atmospheric neutrino reconstruction





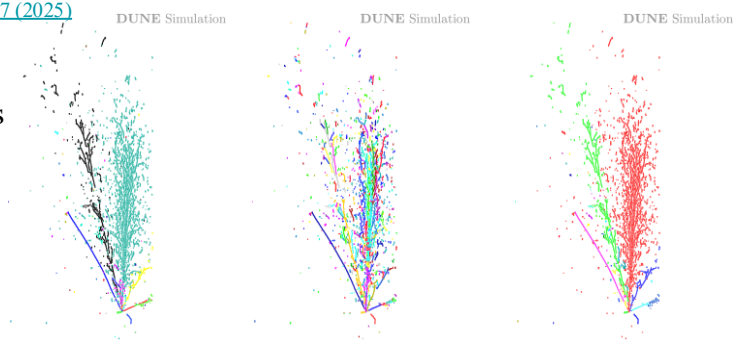


DUNE Collaboration, Reconstruction of atmospheric neutrinos in DUNE's horizontal-drift far-detector module, arXiv: 2601.05697 [hep-ex] (2026)

DUNE Collaboration, Neutrino interaction vertex reconstruction in DUNE with Pandora deep learning, Eur.Phys.J.C 85, 697 (2025) Eur. Phys. J. C 85, 697 (2025)

State-of-the-art shower reconstruction with transformers

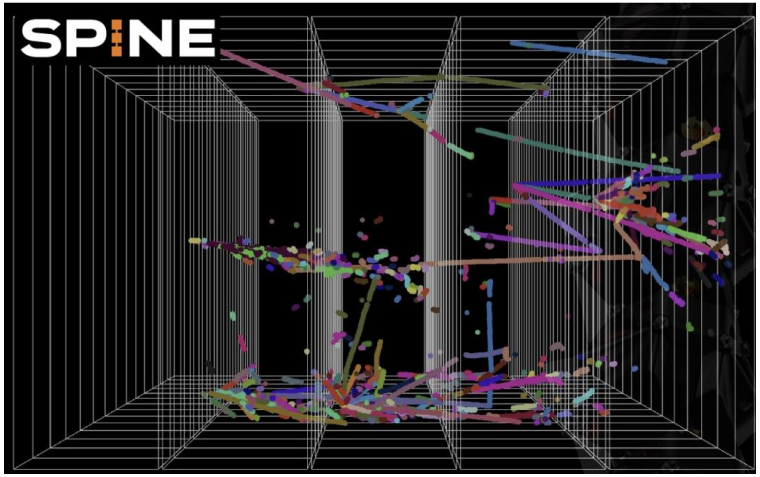
- Significant improvement in dense environments with overlapping showers
- Predicts the similarity matrix between initial simple clusters
- Non-exhaustive list with many other exciting developments underway!



Truth | Simple clustering | Transformer shower growing

Reconstruction Framework: SPINE

- GNNs naturally handle sparse detector hits without requiring a regular grid, while explicitly modeling relationships between hits for particle clustering
- SPINE is an end-to-end neutrino reconstruction framework that combines sparse CNNs and GNNs (including GrapPA) for semantic segmentation, particle clustering, PID, and interaction reconstruction, with applications in ICARUS, SBND, and DUNE (including ProtoDUNE and ND-LAr)



SPINE particle clustering performance

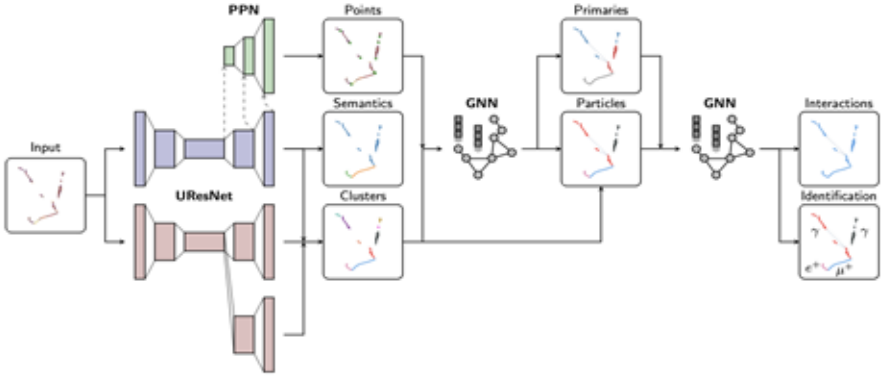
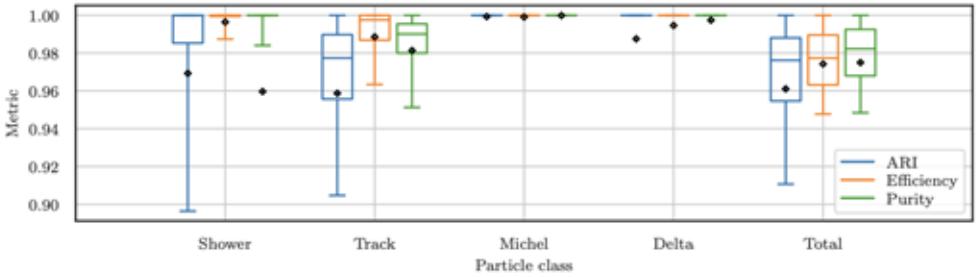


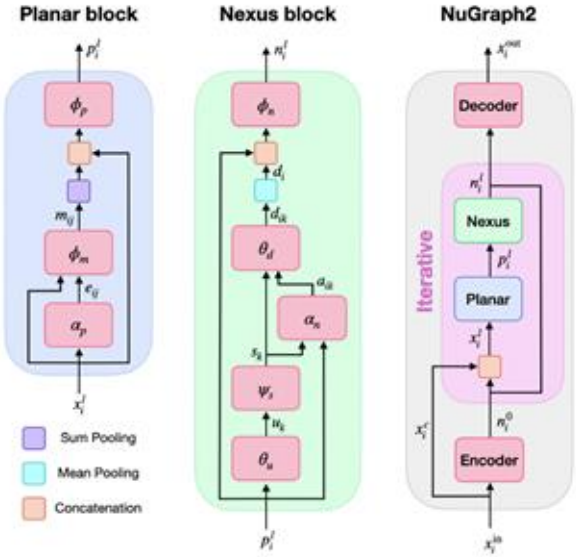
Figure 1: Schematic architecture of the end-to-end, ML-based reconstruction chain for LArTPCs.

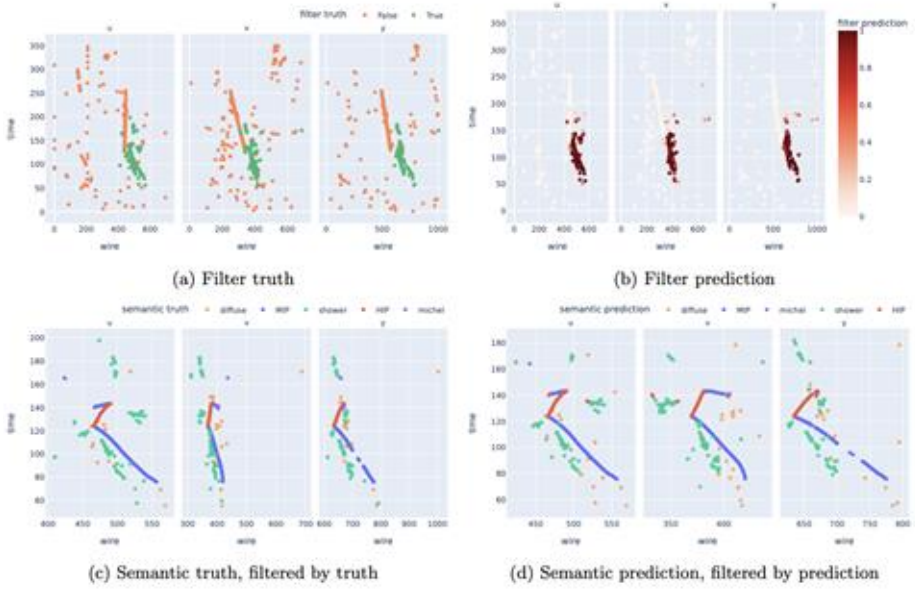


F.Drielsma, K.Terao, L.Domin and D.H.Koh, NeurIPS proceedings, arXiv:2102.01033

Reconstruction Framework: NuGraph

- NuGraph is a multi-task GNN-based neutrino reconstruction framework that performs hit-level semantic segmentation, background rejection, particle clustering, event classification, and vertex regression, with its hierarchical NuGraph3 architecture extending reconstruction from detector hits to particle and interaction-level objects. Applications across MicroBooNE, DUNE, ICARUS, and SBND (particularly through NuGraph3).

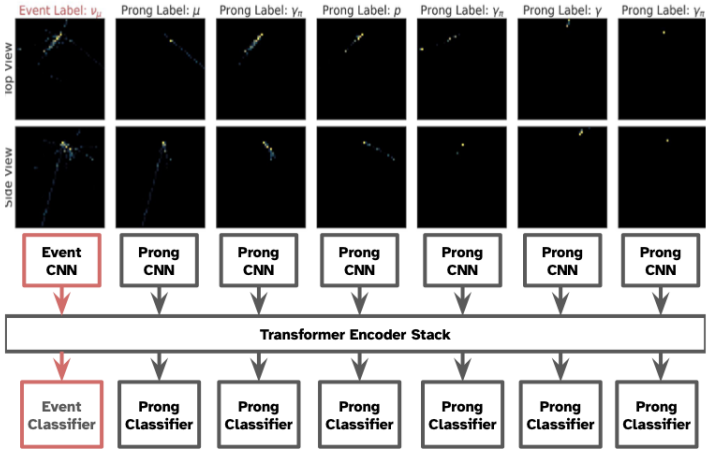




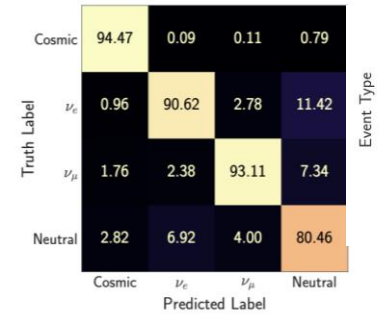
A.Aurisano, V.Hewes, G.Cerati, J.Kowalkowski, C.S.Lee, W.Liao, D.Grzenda, K.Gumpula and X.Zhang, Phys. Rev. D 110, no.3, 3 (2024)

Transformers: Event & Particle Classification

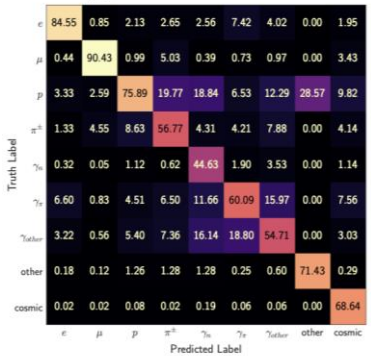
- Transformers enable flexible attention-based modeling of global dependencies and cross-particle interactions, facilitating multi-view information fusion, joint reconstruction tasks, and the development of transferable representations for foundation models.
- Interpretability of network decisions using pixel gradients (saliency) and attention scores



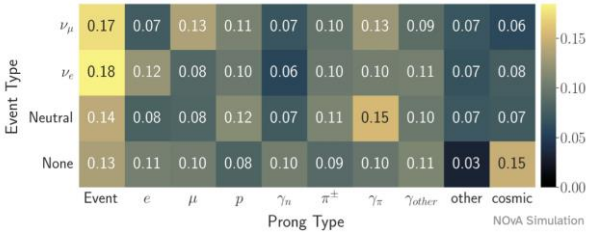
Interpretable Joint Event-Particle Reconstruction with Sparse CNNs and Transformers (2023) (NOvA). arXiv

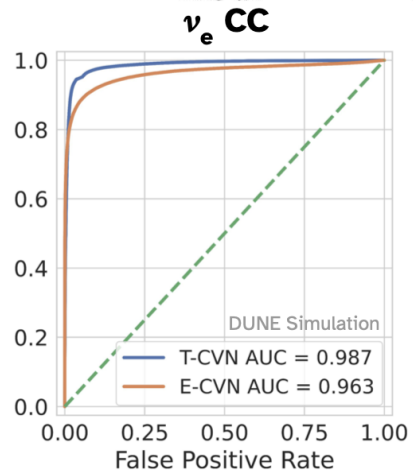


(a) Event Reconstruction



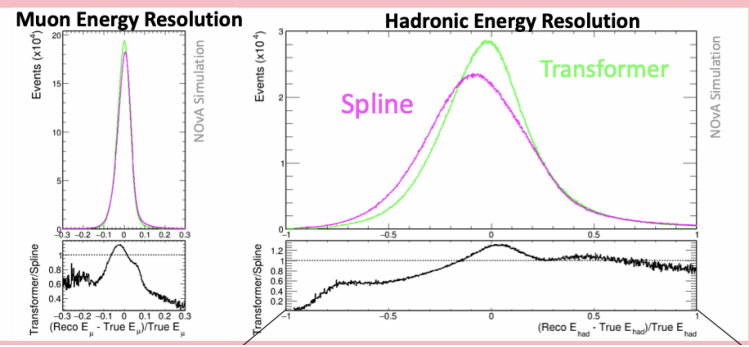
(b) Prong Reconstruction

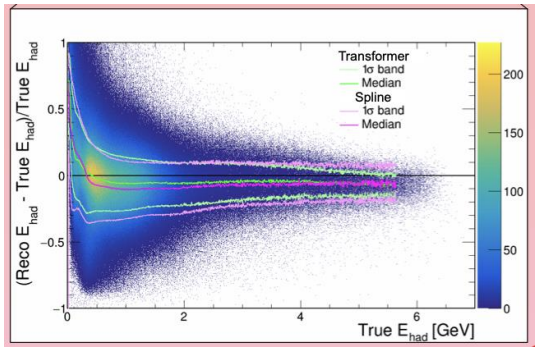




Sparse Convolution Transformers for DUNE FD Event and Particle Classification (2024). Conference

Transformers: Kinematic Regression





Tong, A New Multivariate Transformer Energy Estimator for NOvA (2026) — Conference

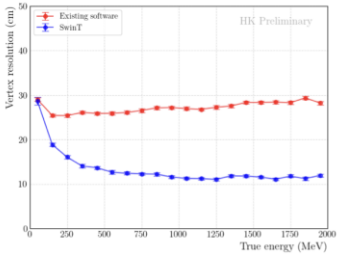


Figure: Vertex res. vs. energy

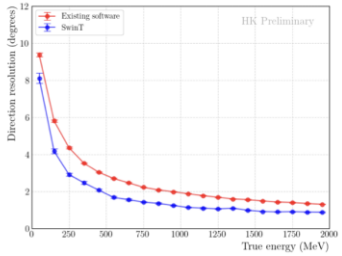


Figure: Direction res. vs. energy

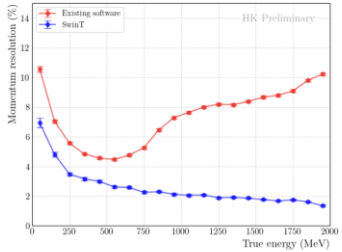
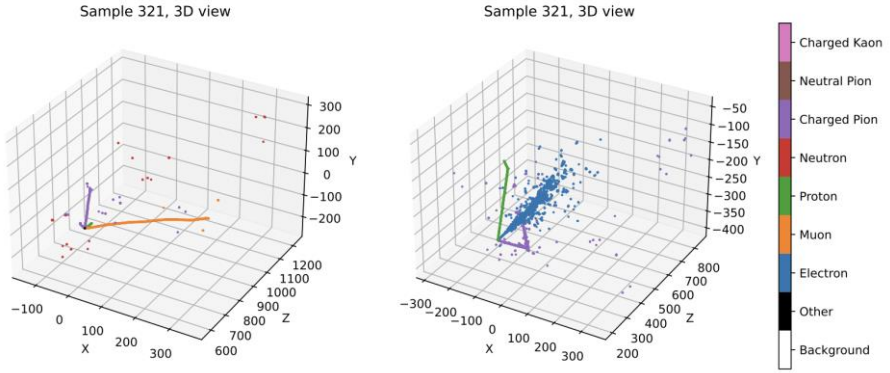


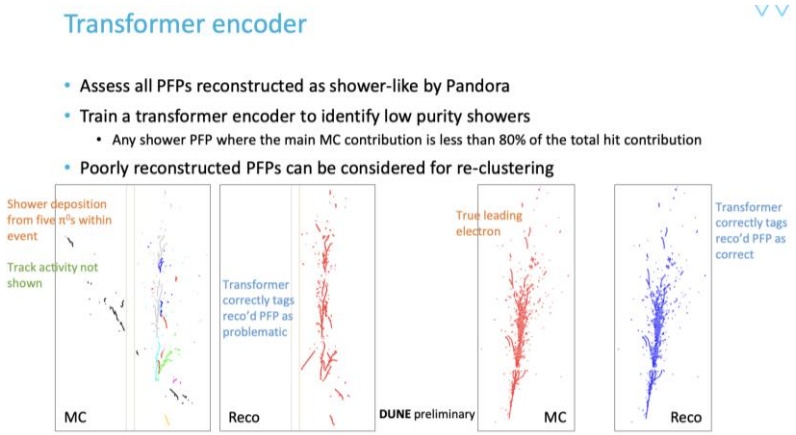
Figure: Momentum res. vs. energy

Chen, Vision Transformers for Event Reconstruction in Water Cherenkov Detector (2025, Hyper-K) —Conference

Transformers: Clustering

- Attention-based modeling of hit relationships and long-range correlations.
- Hierarchical representations capture both local particle structures and global event context.



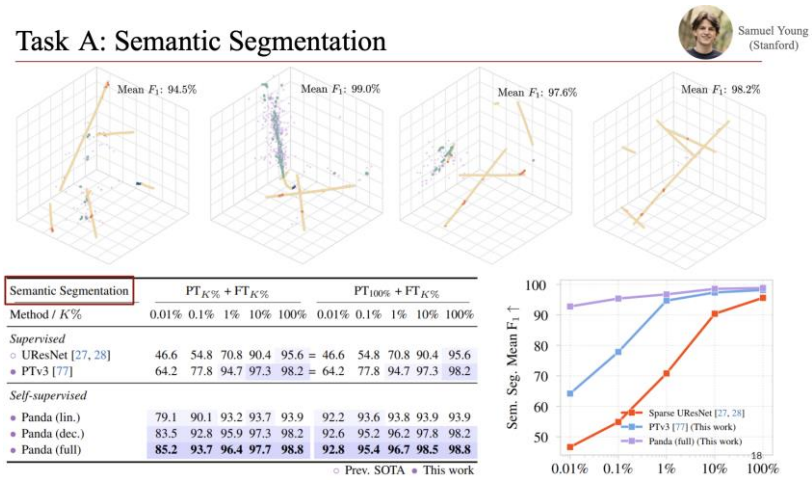


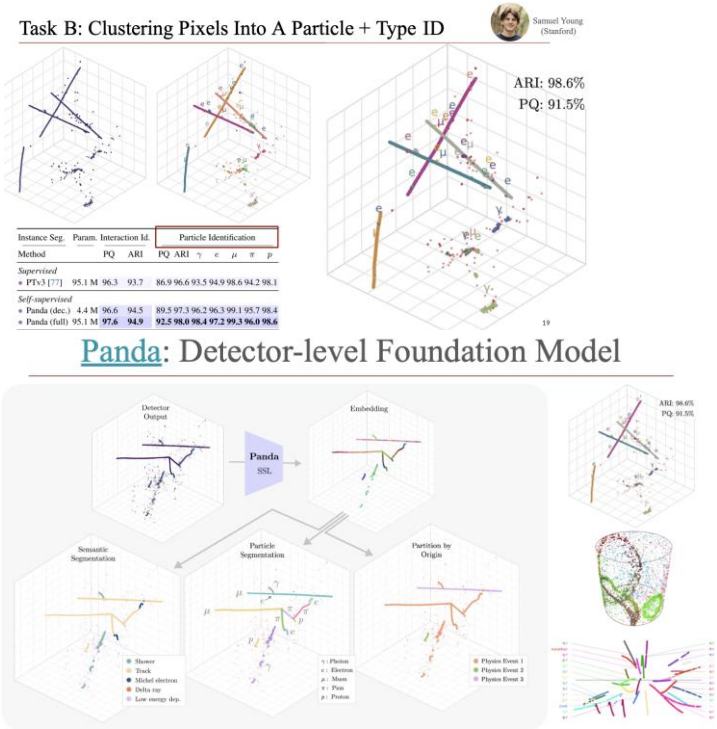
Robles et al., [Particle Hit Clustering and Identification Using Point Set Transformers in Liquid Argon Time Projection Chambers, arXiv:2504.08182 (2025, LArTPC)]

Chappell, [Machine Learning in Pandora (2024)] — Transformer shower-quality classification

Foundation Models

- Foundation models learn general-purpose representations from large-scale datasets, often thru self-supervised pretraining.
- In neutrino reconstruction, they aim to enable knowledge transfer across tasks, detector configurations, and potentially different experiments.





The first step toward sensor-level FMs in HEP

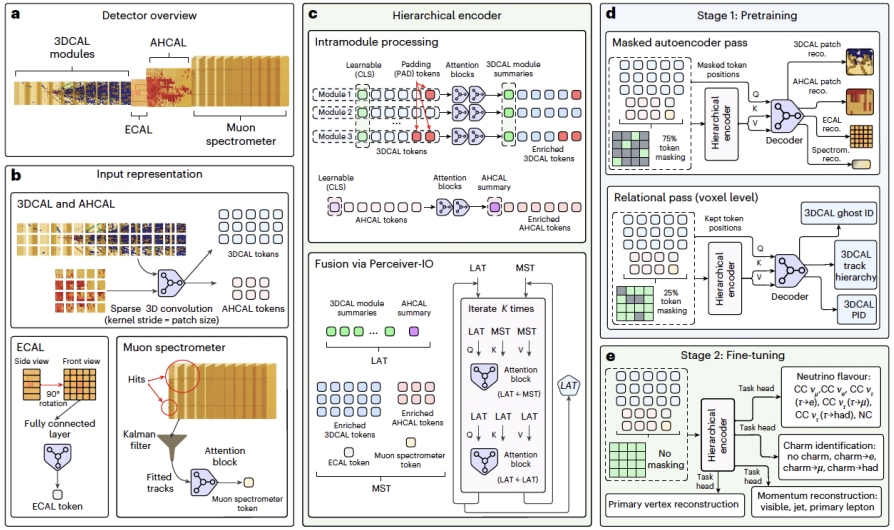
- Mask-based self-supervision: PoLAr-MAE
- Autoregressive token modeling: FM4NPP

Efforts so far focus on foundation across tasks.

Interest/effort in cross-experimental FMs.

Young & Terao, Panda: Self-distillation of Reusable Sensor-level Representations (2025), arXiv:2512.01324

Foundation Models



FASER Cal

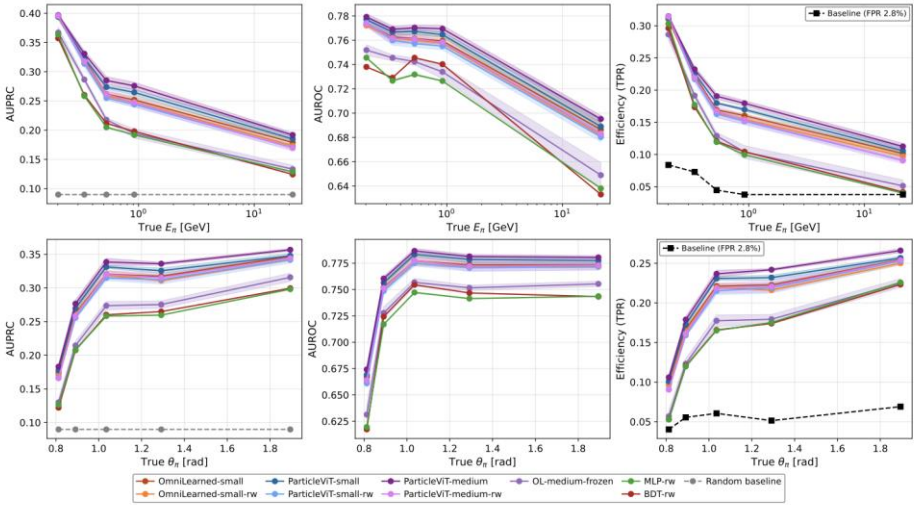


FIG. 2. Discrimination metrics for the CC1$\pi^{\pm}$ tagging task as a function of true pion kinematics: area under the precision–recall curve (AUPRC, left column), area under the ROC curve (AUROC, center), and true positive rate at a fixed false positive rate (right column). The top row uses bins in true pion energy; the bottom row uses bins in true pion polar angle. In each bin, positives are true signal events falling in that bin, while the negative sample is the *full* background, held fixed across bins so that only the in-bin signal population changes. TPR is computed at the baseline false-positive rate matched globally on the full test set; see Fig. 5 for the per-bin matched comparison.

Alonso-Monsalve et al., Towards Foundation-style Models for Energy-frontier Heterogeneous Neutrino Detectors, Nature Machine Intelligence (2026)

Krzmanc et al., Cross-Domain Transfer with Particle Physics Foundation Models (2026, MINERvA Data), arXiv:2604.12364

Multimodal AI

Vision-Language Models (VLMs) enable natural-language interpretation of task-specific neural network predictions and interactive communication with physicists.

Sagar et al., [Adapting Vision-Language Models for Neutrino Event Classification, Commun. Phys. 9, 186 (2026, LArTPC)]

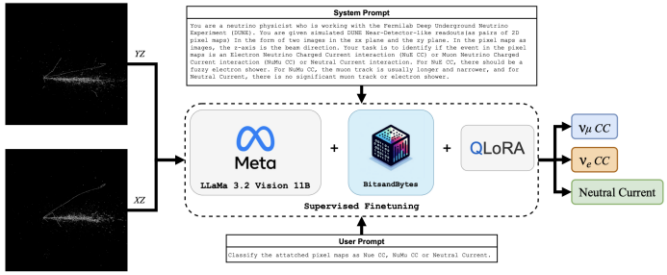


Figure 1: LLaMa 3.2 Vision finetuning pipeline.

Performance

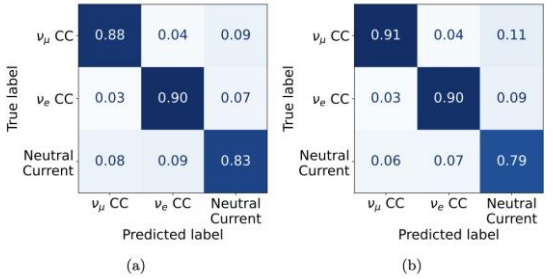


Figure 5: Finetuned LLaMa 3.2 Vision's (a) recall matrix (truth normalized) and (b) precision matrix (prediction normalized).

- ViT-level classification performance
- Robustness to different detector configuration, foundation model potential
- Natural-language explanations
- Human–AI interaction for event analysis

Interpretability

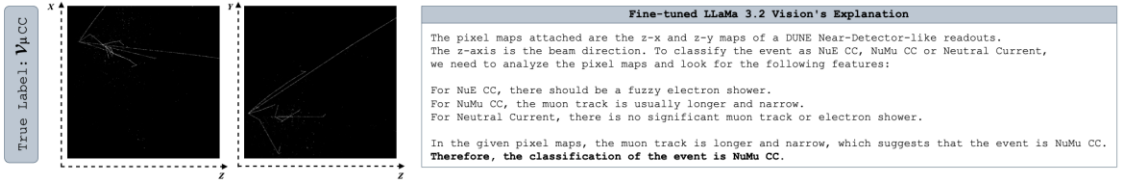


Figure 3: LLaMa 3.2 Vision Prediction Explanation.

Generalization

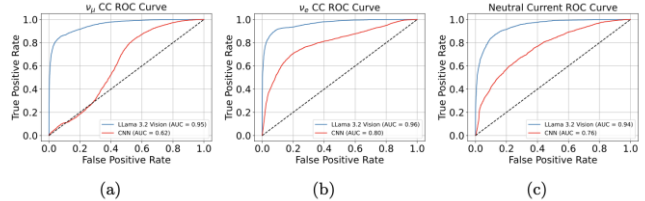
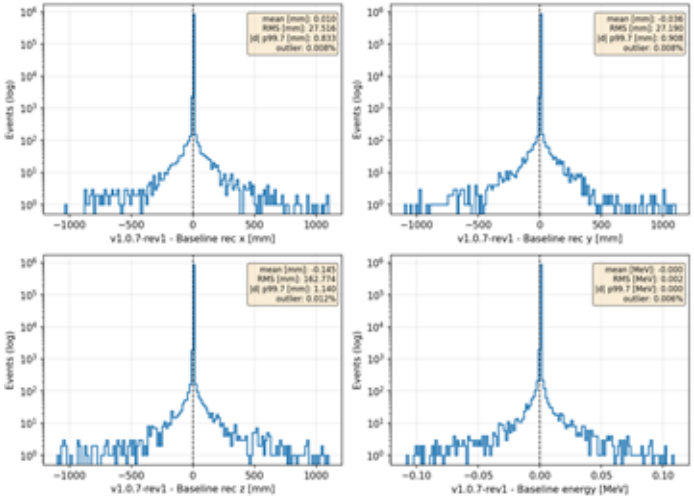


Figure 10: ROC curves for each class (a) ν_μ CC, (b) ν_e CC, and (c) NC comparing performance between the finetuned LLaMa 3.2 Vision and the CNN for generalization testing.

Agentic AI Reconstruction Optimization

Automates multi-step reconstruction optimization and physics analysis through planning, tool execution and validation.



	(A) In-house agent orchestration	(B) Mature agent harness + SKILLs
Reasoning	local LLM (Qwen3.5-4B) + action maps;	frontier model (Claude Opus 4.8) + harness
Retrieval	hybrid RAG, wins at large corpus	curated markdown, lazy-loaded into 1M context
Data access	MCP tools in a fixed pipeline	agentic file/db free navigation
Plots	not multimodal	Can read and analyze images
Deployment	fully offline, $0 per run	cloud API, ~$1–2 /diagnosis, or subscription
Engineering	full-stack system (FastAPI, Vue, Qdrant, Ollama, LangGraph)	~100 lines of markdown, needs pre-built infrastructure (QC database backend/tools)
Choose when	data can't leave site / zero cost per run / can fund the full stack R&D	frontier model permitted / cost per run permitted / engineering time is the scarce resource

Ding, Application of the Dr.Sai Agentic Scientific Workflow in JUNO (NPML 2026) — Vertex / Energy Reconstruction Optimization

Summary

- **Task-specific AI:** CNNs, GNNs, and Transformers address a broad range of neutrino reconstruction tasks, including PID, clustering, and kinematic reconstruction.
- **AI-enhanced reconstruction pipelines:** Frameworks such as Wire-Cell, Pandora, SPINE, and NuGraph increasingly integrate AI/ML into particle reconstruction.
- **Foundation Models:** Large-scale pretraining enables the development of transferable representations across reconstruction tasks, detector configurations, and potentially different experiments.
- **Multimodal and Agentic AI:** Emerging approaches enable natural-language interaction, interpretable reconstruction, and automated optimization of reconstruction workflows.

Thank you! 감사합니다!

Backup

Representative CNN-Based Neutrino Reconstruction Methods

Reconstruction Task	Experiment	Representative Paper & Reference
Hit / Semantic Segmentation	MicroBooNE	MicroBooNE Collab., Semantic Segmentation with a Sparse CNN for Event Reconstruction in MicroBooNE, Phys. Rev. D 103, 052012 (2021)
	DUNE	DUNE Collab., Separation of Track- and Shower-like Energy Deposits in ProtoDUNE-SP Using a CNN (2022), EPJC
Vertex Reconstruction	DUNE	DUNE Collab., Neutrino Interaction Vertex Reconstruction in DUNE with Pandora Deep Learning, EPJC (2025)
	MINERvA	Song et al., Deep Learning for Vertex Reconstruction of Neutrino-Nucleus Interaction Events, ICASSP (2019)
	JUNO (Reactor)	Qian et al., Vertex and Energy Reconstruction in JUNO with Machine Learning Methods, NIM A 1010, 165527 (2021)
Track Reconstruction	JUNO (Reactor)	Liu et al., Muon Reconstruction with a CNN in the JUNO Detector, Radiat. Detect. Technol. Methods 5, 364-372 (2021)
Particle Identification (PID)	NOvA	Psihas et al., Context-Enriched Identification of Particles with a Convolutional Network, Phys. Rev. D 100, 073005 (2019)
	MicroBooNE	MicroBooNE Collab., CNN for Multiple Particle Identification, Phys. Rev. D 103, 092003 (2021)
	T2K	Arai, CNN-Based Particle Identification in SuperFGD (2025). Conference Report
	Hyper-Kamiokande	Atta et al., Enhancing Event Reconstruction in Hyper-Kamiokande with Machine Learning: A ResNet Implementation (2026), arXiv:2604.13503
Event Classification / Selection	NOvA	Aurisano et al., A Convolutional Neural Network Neutrino Event Classifier, JINST 11, P09001 (2016)
	Daya Bay (Reactor)	Racah et al., Exploring Raw HEP Data Using Deep Neural Networks, PoS(ICHEP2016) 889
Kinematic Regression (Energy / Vertex / Momentum / Direction)	NOvA	Baldi et al., Improved Energy Reconstruction in NOvA with Regression CNNs, Phys. Rev. D 99, 012011 (2019)
	JUNO (Reactor)	Qian et al., Vertex and Energy Reconstruction in JUNO with Machine Learning Methods, NIM A 1010, 165527 (2021)
	Hyper-Kamiokande	Atta et al., ResNet Implementation (2026), arXiv:2604.13503

Representative GNN-Based Neutrino Reconstruction Methods

Reconstruction Task	Experiment	Representative Paper & Reference
Hit Clustering / Instance Segmentation	Generic LArTPC	Drielsma et al., Clustering of Electromagnetic Showers and Particle Interactions with GNNs, Phys. Rev. D 104, 072004 (2021)
	ICARUS	ICARUS Collab., First Search for Sterile Neutrino Oscillation at ICARUS (2026), Appendix B — SPINE, arXiv:2603.22557
Hit Classification / Background Rejection	MicroBooNE	Hewes et al., Graph Neural Network for Neutrino Physics Event Reconstruction, Phys. Rev. D 110, 032008 (2024)
	T2K	Alonso-Monsalve et al., GNN for 3D Classification of Ambiguities and Optical Crosstalk, Phys. Rev. D 103, 032005 (2021)
Particle Association / Track Matching	DUNE	Micallef, Track Matching during Reconstruction Using GNNs Across DUNE's Near Detectors (2026) — Conference
Interaction Reconstruction / Event Classification	ICARUS	ICARUS Collab., SPINE / GrapPA Interaction Clustering (2026), arXiv:2603.22557
	Hyper-Kamiokande	Le Blévec, Graph Neural Networks for Hyper-K Reconstruction (2025) — Conference
Kinematic Regression (Energy / Vertex / Direction)	JUNO (Reactor)	Qian et al., Vertex and Energy Reconstruction in JUNO with Machine Learning Methods, NIM A 1010, 165527 (2021) — DeepSphere GNN for energy
	TAO (Reactor)	Hariharan, Vertex Reconstruction in JUNO-TAO Using Deep Learning (2022) — GNN, Conference
	JUNO (Reactor)	Li et al., A Multi-purposed Reconstruction Method Based on Machine Learning for Atmospheric Neutrino at JUNO, J. Phys.: Conf. Ser. 3206, 012090 (2026)

Representative Transformer-Based Neutrino Reconstruction Methods

Reconstruction Task	Experiment	Representative Paper & Reference
Hit Clustering / Segmentation	NOvA	Robles et al., Heterogeneous Point Set Transformers for Segmentation of Multiple View Particle Detectors (2025), arXiv:2510.09659
	Generic LArTPC	Robles et al., Particle Hit Clustering and Identification Using Point Set Transformers (2025), arXiv:2504.08182
	MicroBooNE	Da Silva, LoRA Fine-Tuning on Sonata Pre-trained LArTPC Point Cloud Network (2026) — Conference
Particle Identification (PID)	NOvA	Shmakov et al., Interpretable Joint Event-Particle Reconstruction with Sparse CNNs and Transformers (2023), arXiv:2303.06201
	DUNE	Yankelevich & Shmakov, Sparse Convolution Transformers for DUNE FD Event and Particle Classification (2024) — Conference
	T2K	Kiseeva, Transformer-Based Approaches to Particle Identification in SuperFGD (2025) — Conference
Event Classification	NOvA	Shmakov et al., Interpretable Joint Event-Particle Reconstruction (2023), arXiv:2303.06201
	DUNE-related LArTPC	Moretti et al., Assessment of Few-Hits Machine Learning Classification Algorithms, EPJ Plus 139, 723 (2024)
	MINERvA	Krzmanc et al., Cross-Domain Transfer with Particle Physics Foundation Models (2026), arXiv:2604.12364
	JUNO (Reactor)	Fang et al., Deep Learning-Based ^{14}C Pile-Up Identification in JUNO (2026), arXiv:2603.01419
Kinematic Regression (Energy / Vertex / Momentum / Direction)	NOvA	Tong, A New Multivariate Transformer Energy Estimator for NOvA (2026) — Conference
	MINERvA	Krzmanc et al., Cross-Domain Transfer with Particle Physics Foundation Models (2026), arXiv:2604.12364
	JUNO (Reactor)	Li et al., A Multi-purposed Reconstruction Method Based on Machine Learning for Atmospheric Neutrino at JUNO, J. Phys.: Conf. Ser. 3206, 012090 (2026)
Optical / Timing Reconstruction	SBND	Crespo Anadón et al., Neural Network-Based Reconstruction Using the SBND Photon Detection System (2026) — Technical Report
	DUNE-related LArTPC	Towards Self-Supervised Optical Signal Reconstruction in LArTPCs (2026) — Conformer, Conference
Reconstruction Quality Assessment	DUNE	Chappell, Machine Learning in Pandora (2024) — Transformer shower-quality classification
Event-Level Reconstruction	Hyper-Kamiokande	Chen, Vision Transformers for Event Reconstruction in Water Cherenkov Detector (2025) — Conference
Multi-task Reconstruction / Pretraining	FASERCal (Proposed)	Alonso-Monsalve et al., Towards Foundation-Style Models for Energy-Frontier Heterogeneous Neutrino Detectors, Nature Machine Intelligence (2026)

Representative Foundation Models in Neutrino Experiments

Reconstruction Task	Experiment / Dataset	Representative Paper & Reference
Particle Clustering / Instance Segmentation	Generic LArTPC	Young & Terao, Panda: Self-distillation of Reusable Sensor-level Representations (2025), arXiv:2512.01324
	Multi-detector Simulation	Young et al., Panda Diplomacy (2026), arXiv:2609.00611
Semantic Segmentation / PID	Generic LArTPC	Young et al., Particle Trajectory Representation Learning with Masked Point Modeling, Mach. Learn.: Sci. Technol. 7, 025023 (2026)
	MicroBooNE	Da Silva, LoRA Fine-Tuning on Sonata Pre-trained LArTPC Point Cloud Network (NPML 2026)
	FASERCal (Proposed)	Alonso-Monsalve et al., Towards Foundation-style Models for Energy-frontier Heterogeneous Neutrino Detectors, Nature Machine Intelligence (2026)
Event Classification	MINERvA	Krzmanc et al., Cross-Domain Transfer with Particle Physics Foundation Models (2026), arXiv:2604.12364
Kinematic Regression	MINERvA	Krzmanc et al., Cross-Domain Transfer with Particle Physics Foundation Models (2026), arXiv:2604.12364
	FASERCal (Proposed)	Alonso-Monsalve et al., Towards Foundation-style Models for Energy-frontier Heterogeneous Neutrino Detectors, Nature Machine Intelligence (2026)
Optical / Timing Reconstruction	DUNE ND (R&D)	Smith, Towards Self-Supervised Optical Signal Reconstruction in LArTPCs (NPML 2026)
Self-Supervised Representation Learning	Wire-Cell / LArTPC	Vicenzi, Development of a Self-Supervised Foundation Model for Wire-Cell (NPML 2026)