Fermilab

Blueprint Workshop: Towards a National-Scale AI Collaboration in HEP, May 18th, 2026

Neutrinos

(mostly from accelerators)

Alex Himmel
Senior Scientist, Fermilab

U.S. DEPARTMENT of ENERGY
Fermi National Accelerator Laboratory is managed by
FermiForward for the U.S. Department of Energy Office of Science

Major Goals in Accelerator Neutrino Physics

This is what we are trying to achieve, and hope AI can help

- Neutrino oscillations
- Neutrino cross sections
  - Both to support oscillations and in their own right
- Searches for BSM physics
  - Non-standard oscillations (sterile neutrinos, Lorentz violation, Nnn-standard interactions)
  - Coincident production (dark matter, heavy neutral leptons, etc.)
- Astroparticle physics
  - Galactic and extragalactic neutrinos
  - Multi-messenger astronomy
  - Supernovas
- Lots more out there. Some things I won't talk much about:
  - Reactor neutrino experiments
  - Direct mass measurements
  - Neurinoless double beta decay

Experiment Landscape

Accelerator and atmospheric experiments

Recently completed but still active:

- µBooNE (BSM, Cross sections)
- MINERvA (Cross sections)

Currently taking data:

- T2K (Oscillations, Cross sections)
- NOvA (Oscillations, BSM, Cross sections)
- SBN: SBND+ICARUS (BSM, Cross sections)
- Super-Kamiokande (Oscillations, BSM, Astroparticle)
- IceCube (Oscillations, BSM, Astroparticle)
- KM3NET (Oscillations, BSM, Astroparticle)

Future:

- DUNE (everything)
- Hyper-Kamiokande (everything)

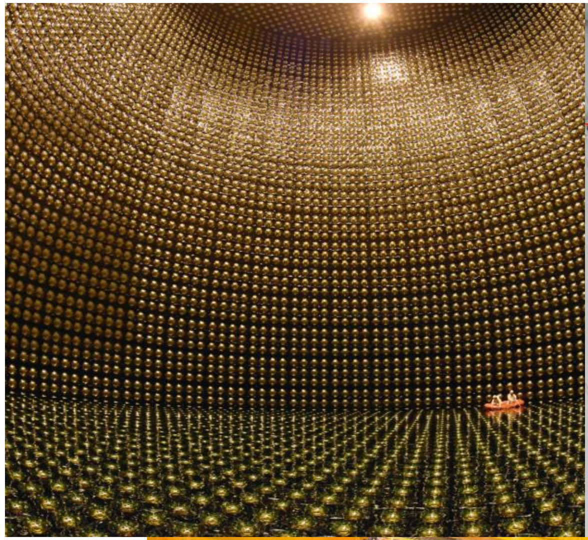

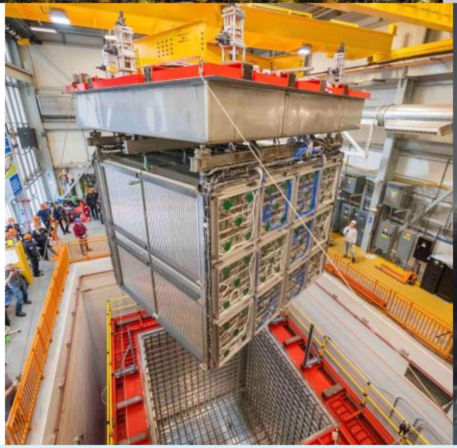
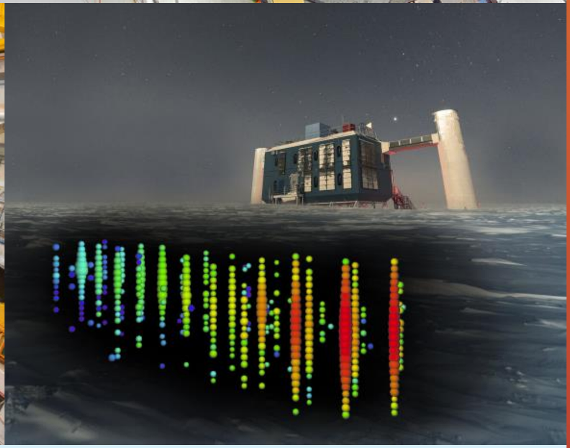
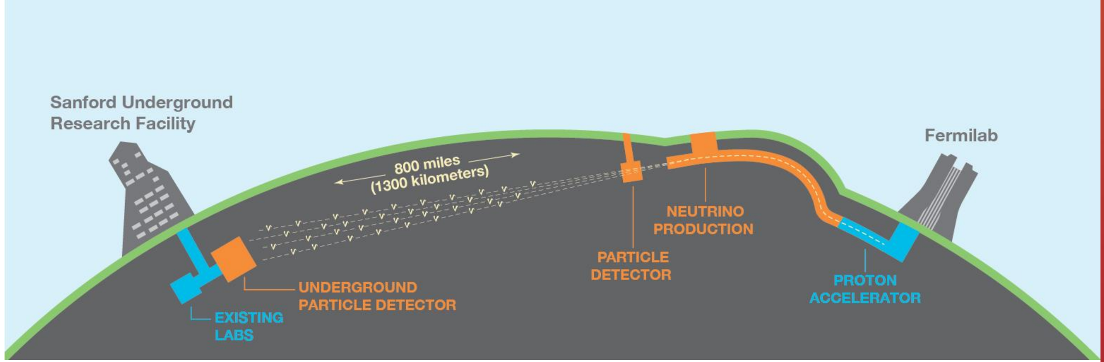

How to Measure Oscillations

Observe flavor change as a function of energy over a long distance while mitigating uncertainties on neutrino flux, cross sections, and detector response.

How to Measure Oscillations

Particle ID

Reconstruction

Observe flavor change as a function of energy over a long distance while mitigating uncertainties on neutrino flux, cross sections, and detector response.

Key systematics – models

Extrapolation and control samples

Current AI Applications: Particle ID

How we're using AI today

- First application of deep learning in neutrinos (2016).
- Neutrino detectors tend to to be large, homogenous, tracking calorimeters
- Well-suited to CNNs and related architectures
  - Challenge has been on adapting to non-natural images (ex: sparsity, granularity, size, etc)
- Generally developed within collaborations, though ideas of course spread around

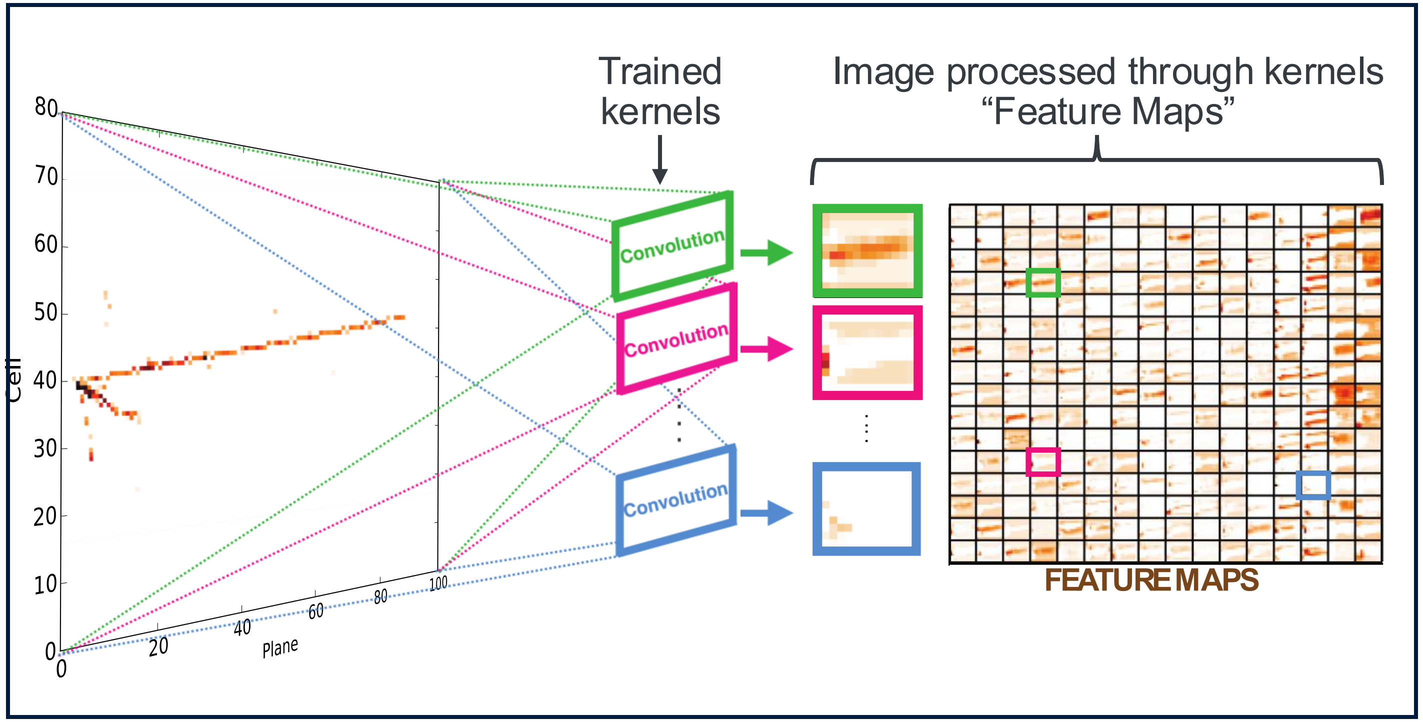

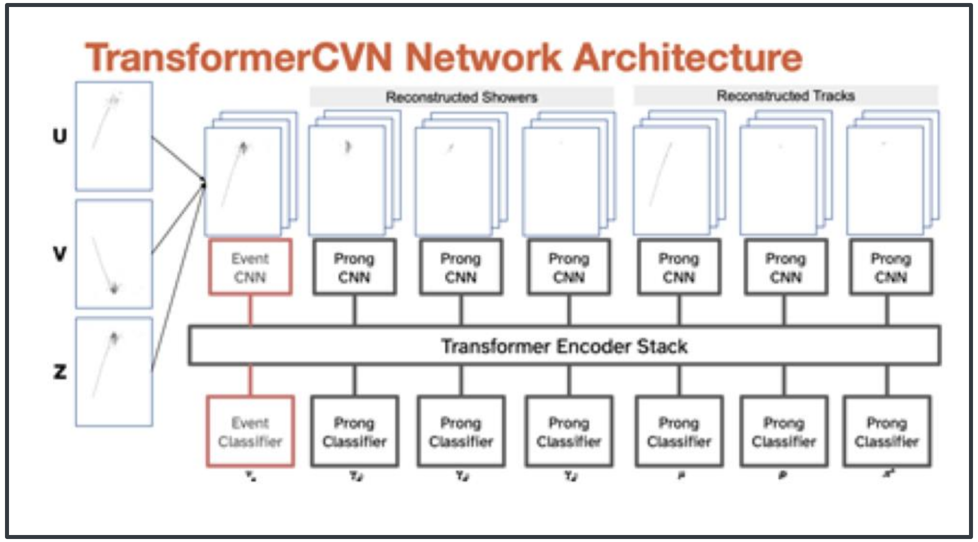

From J. Bian, *NuFact2025*, 9/4/25

Current AI Applications: Reconstruction

How we're using AI today

- Examples: ROI, Clustering, view matching, sub-component identification, regression, etc.
- More diversity: CNNs, GNNs, Transformers, etc.
- Some cross-experiment collaborations here already:

SPINE | Pandora | ExaTrkX | WireCell

H. Yu et al., JINST 16 (2021) 01, P01036

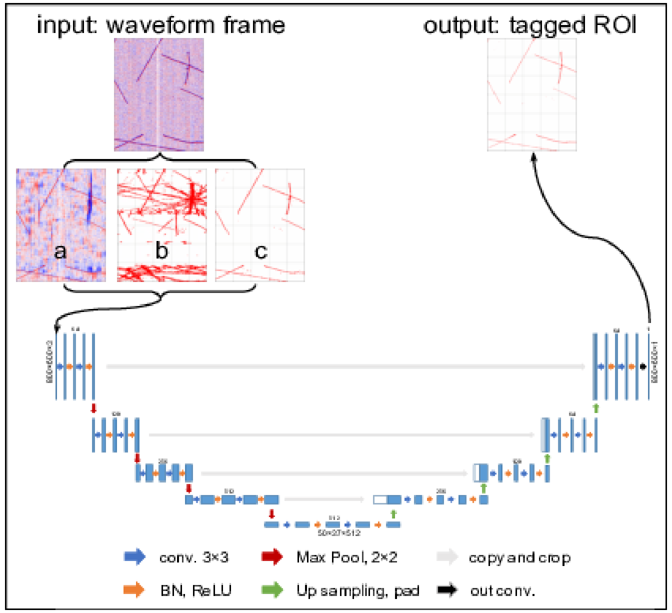

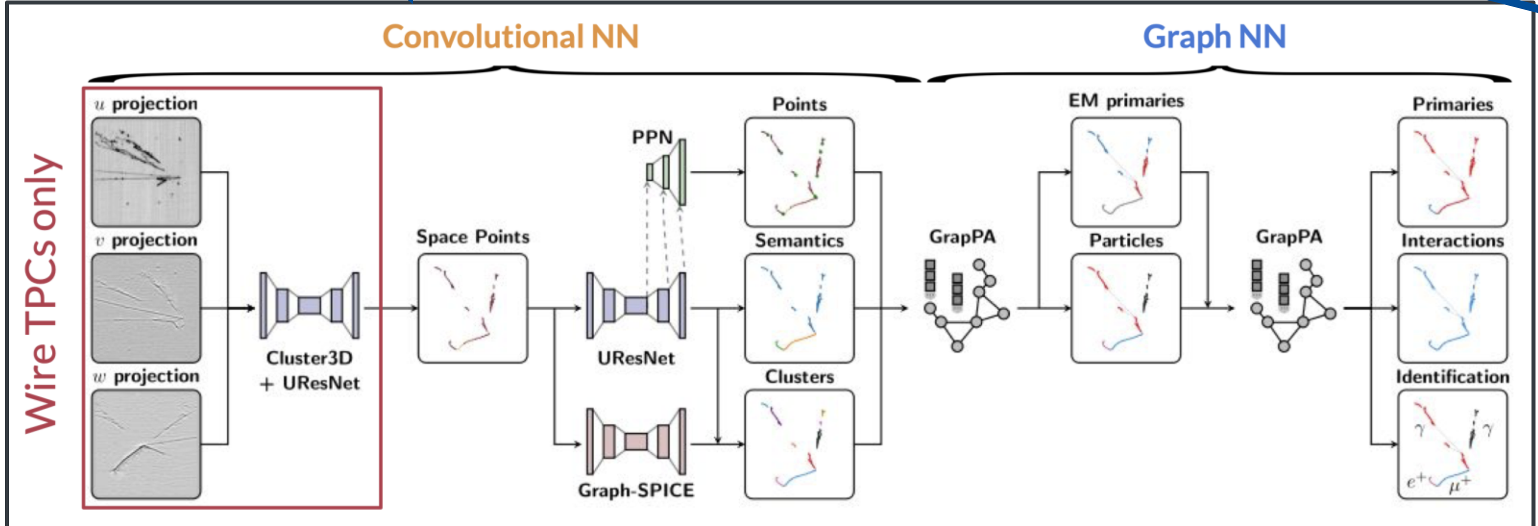

Paper: arXiv:2102.01033

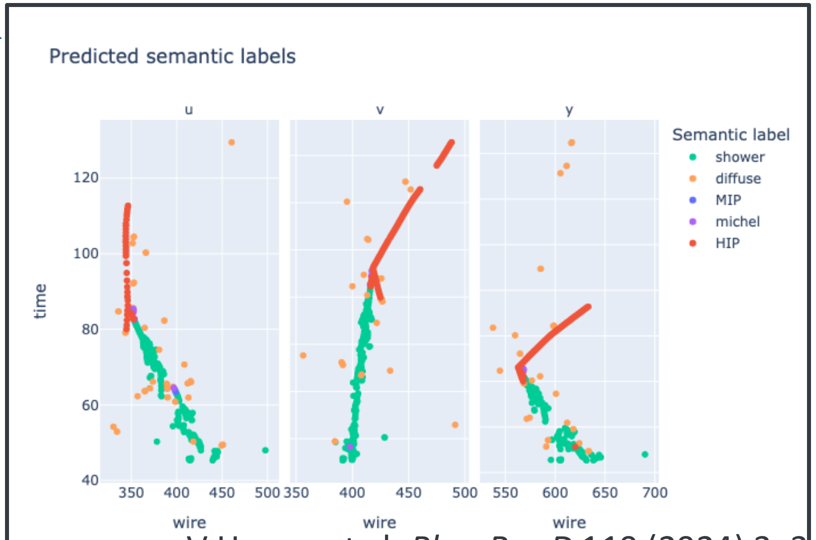

V Hewes et al. *Phys.Rev.D* 110 (2024) 3, 3

Future AI Applications

A *highly* incomplete list, trying to capture the interest of attendees

- Agentic analysis tools
- Agentic detector & computing operations
- Training & On-boarding

New tooling to make doing physics faster and easier

- Learned cross-section models
- Surrogate models for detector simulation
- Foundation and world models (for reco, for interactions, etc.)

New paradigms for theory-experiment collaboration - improve key models

- Realtime/Nearline inference (ex: supernovas)
- Simulation-based inference
- Uncertainty quantification
- Differentiable simulation

Novel approaches to current (computational) problems

Datasets and Organizational Thoughts

- Shared technical challenges: adapting to HPC/GPU, reproducibility, etc.
- A challenge: the neutrino community is fragmented into many smaller experiments.
  - As we centralize into larger collaborations (SBN, DUNE, HK) maybe this will improve
- Also, most neutrino data is not public, with some notable listed exceptions below
  - Partly because it's a lot of work to make it useful. Maybe AI can help here!

Neutrino Open Data

Experiment	(1) Data	(2) Metadata	(3) Analysis Data	(4) Tools
MicroBooNE	2D wire-time images, event data — HDF5, artroot — Zenodo	dataset docs — Fermilab site	reduced datasets — Zenodo	Python data loaders, example Jupyter notebooks — GitHub (OpenSamples)
MINERvA	tabular ntuples — ROOT — Fermilab server	docs — Fermilab site	reco vars + systematics — Fermilab server	ROOT macros (event reading, plotting), systematics weighting scripts — GitHub (MinervaExpt)
Daya Bay	tabular event data — HDF5, NPZ, ROOT — Zenodo	metadata — Zenodo	analysis dataset (IBD + inputs) — Zenodo	Python analysis package (dayabay-model), fit scripts — GitHub
NOvA	histograms, tables — ROOT/plots — Fermilab publicdocs	docs — Fermilab site	oscillation / xsec products — Fermilab publicdocs	
T2K	fit outputs, histograms — ROOT — Zenodo	metadata — Zenodo	fit results — Zenodo	ROOT macros (plot extraction, validation) — Zenodo
IceCube	event-level data (time, charge, position) — HDF5/CSV — IceCube data release site	docs — IceCube site	event catalogs — IceCube site	
PILArNet	2D/3D sparse images, point clouds — numpy/HDF5 — OSF	dataset paper — arXiv	labeled dataset — OSF	dataset interface scripts (data loading, preprocessing) — project repo
NuBench	detector hits (graph/point-cloud, tabular) — Parquet, SQLite — ERDA	docs — arXiv + GitHub	benchmark tasks — GitHub	data loaders, training scripts, model implementations (GraphNeT-based) — GitHub

3

From J. Bian, *Open (and Closed) Data in the Age of AI Workshop, Chicago*, 4/22/26