Overview:
ML in Neutrino Physics

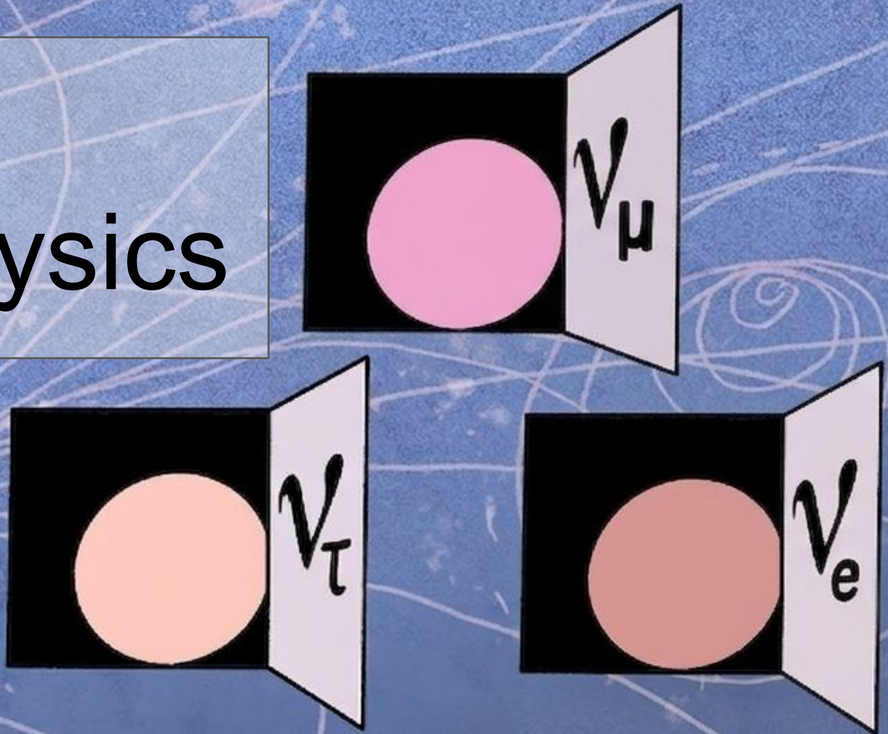

Yifan Chen (SLAC)
ML4FP @ Georgia Tech
2026/06/02

Neutrinos!

- 3 flavours, charge neutral, small masses
- Neutrinos interact rarely
- Neutrino oscillations (neutrino mixing, mass splitting)

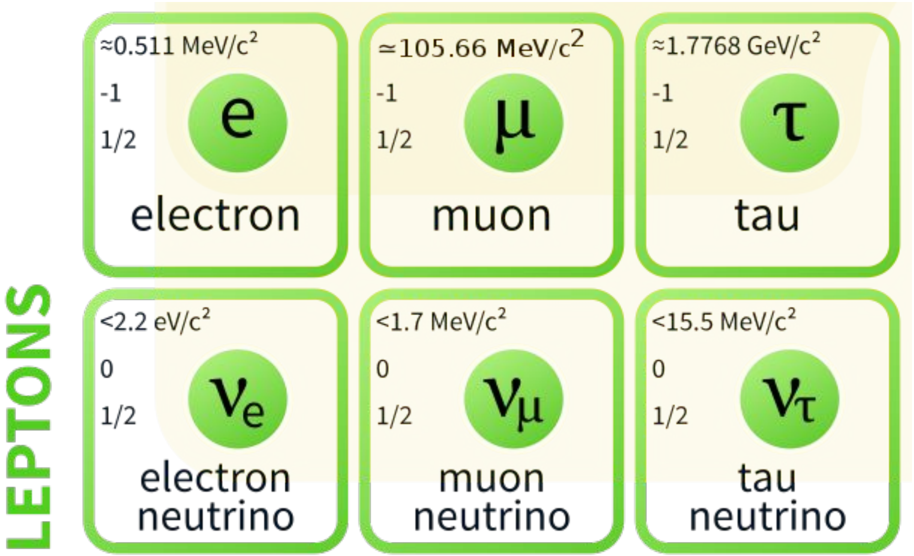

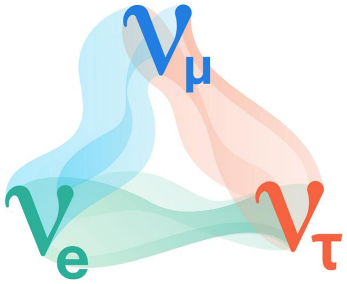

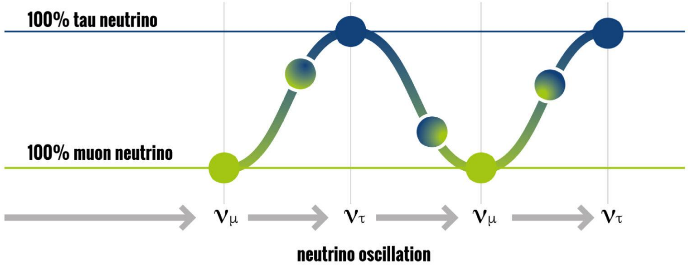

Neutrino sources and energies

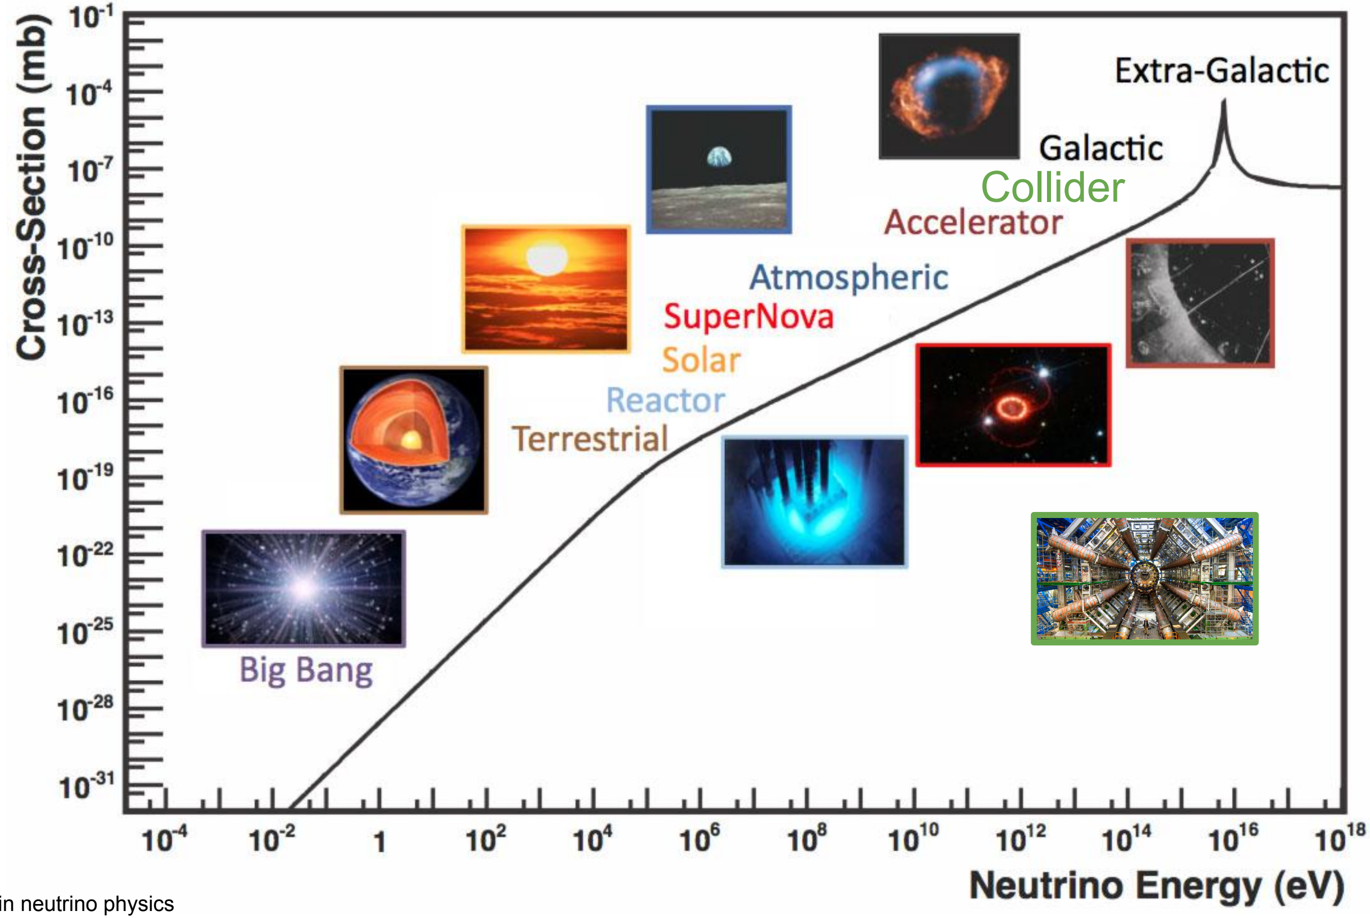

Neutrino related open questions

- Neutrino mass ordering?
- Charge-parity violation in the lepton sector?
- Sterile neutrinos?
- Are neutrinos Majorana or Dirac particles?
- Absolute neutrino mass scale?
- Origin of the neutrino mass?
- Supernova neutrinos?
- Diffuse supernova neutrino background?
- Astrophysical neutrino sources?
- Ultra-high-energy neutrinos?
- Geo-neutrinos?
- Cosmic neutrino background?

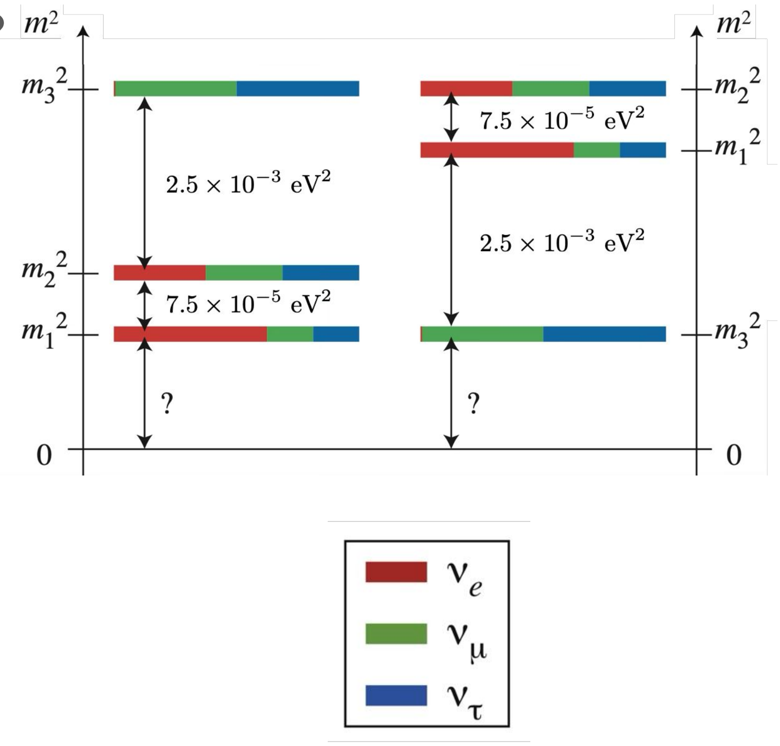

Neutrino experiments

Accelerator neutrinos

DUNE
DEEP UNDERGROUND NEUTRINO EXPERIMENT

Reactor neutrinos

JUNO

Collider neutrinos

FASER

Astrophysical neutrinos

ICECUBE

Neutrinoless double beta decay

LEGEND

Neutrino mass

KATRIN
KARLSRUHE TRITIUM NEUTRINO EXPERIMENT

~~Overview:~~
ML in Neutrino Physics

Tasting Sampler

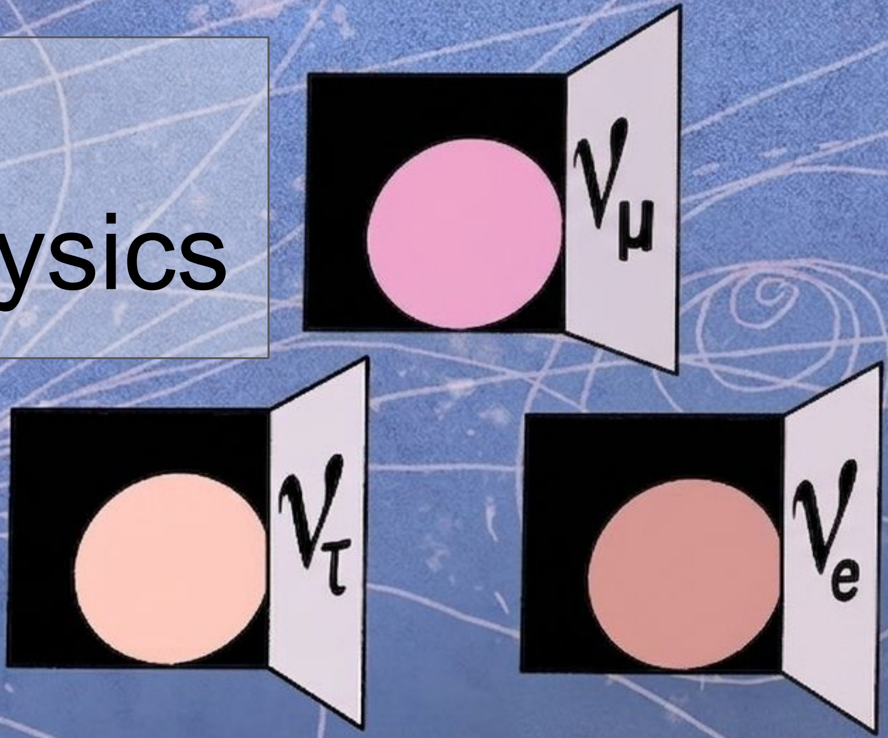

Further discussion at the neutrino parallel session!

Neutrino dataset: Neutrino detector

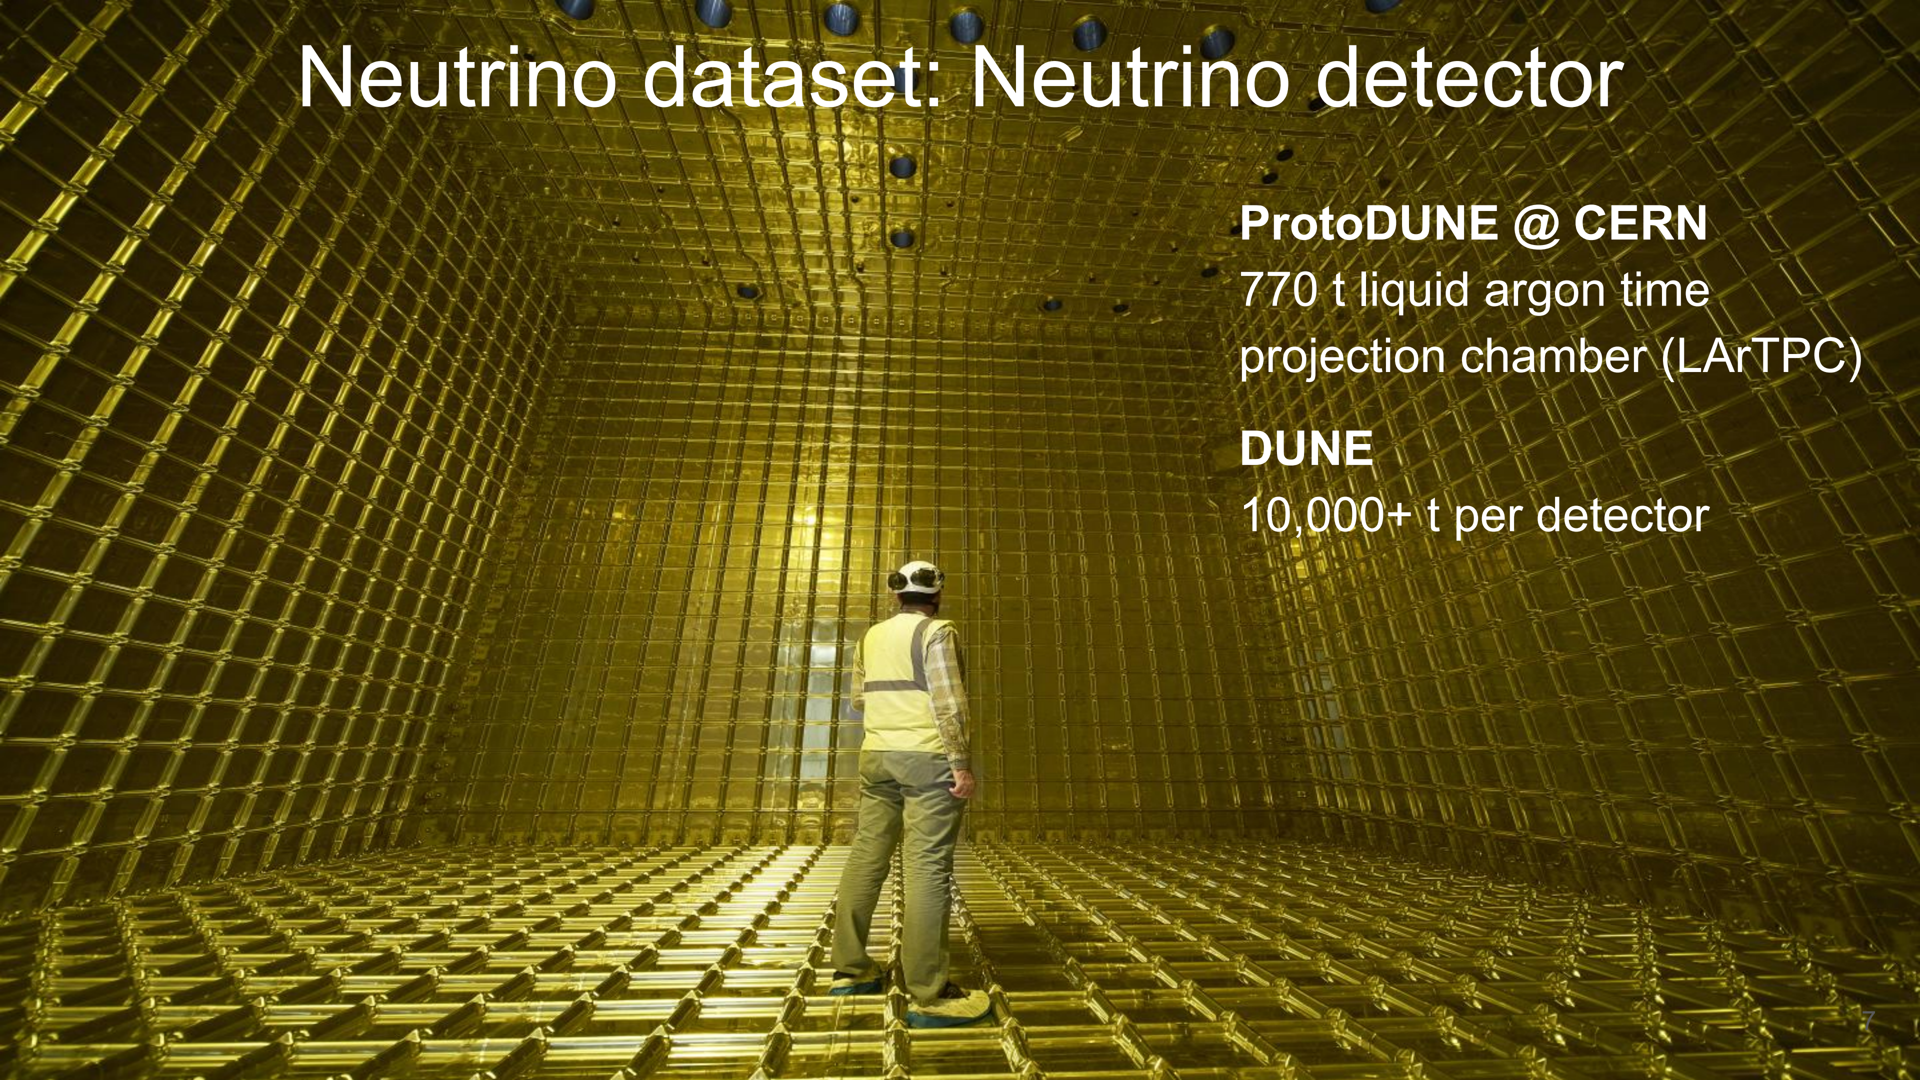

ProtoDUNE @ CERN

770 t liquid argon time projection chamber (LArTPC)

DUNE

10,000+ t per detector

Neutrino dataset: Neutrino detector

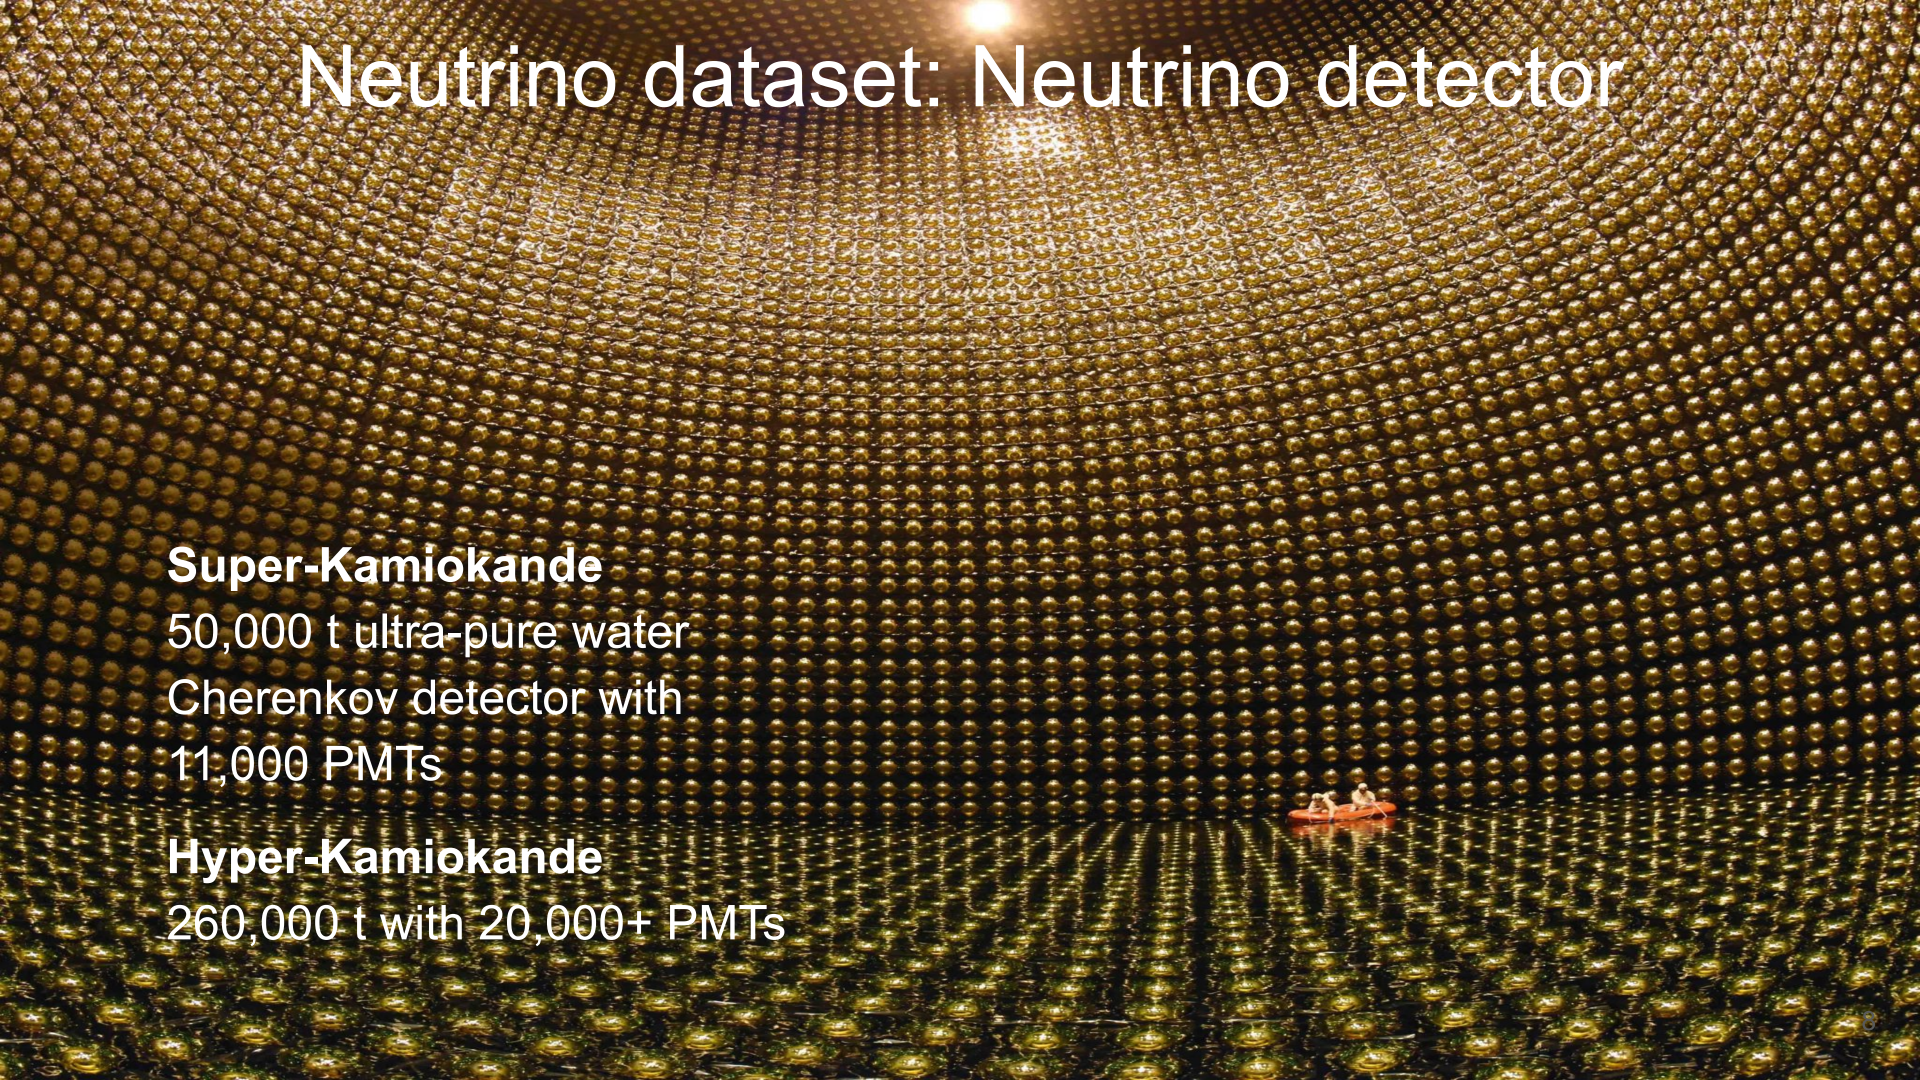

Super-Kamiokande
50,000 t ultra-pure water
Cherenkov detector with
11,000 PMTs

Hyper-Kamiokande
260,000 t with 20,000+ PMTs

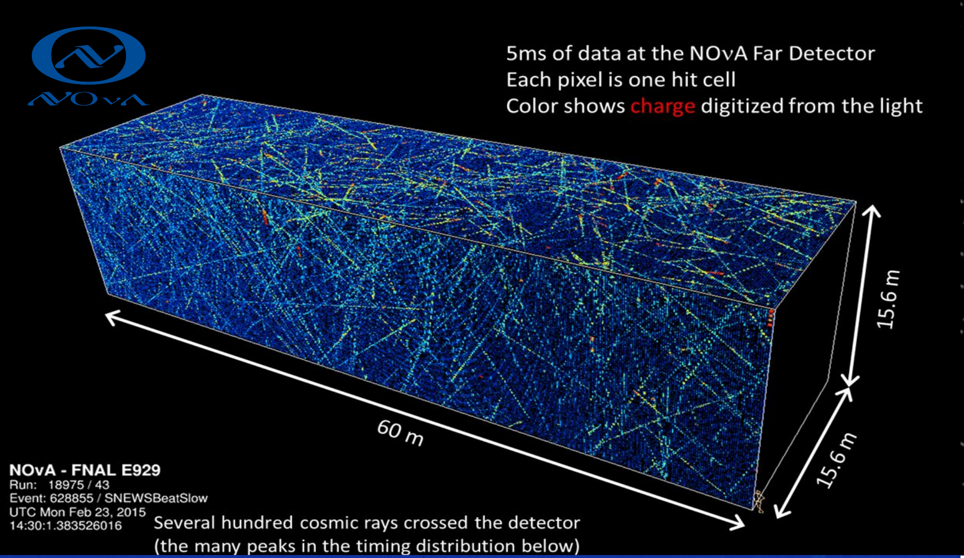
5ms of data at the NOvA Far Detector
Each pixel is one hit cell
Color shows charge digitized from the light
15.6 m
60 m
15.6 m
NOvA - FNAL E929
Run: 18975 / 43
Event: 628855 / SNEWSBeatSlow
UTC Mon Feb 23, 2015
14:30:1.383526016
Several hundred cosmic rays crossed the detector
(the many peaks in the timing distribution below)

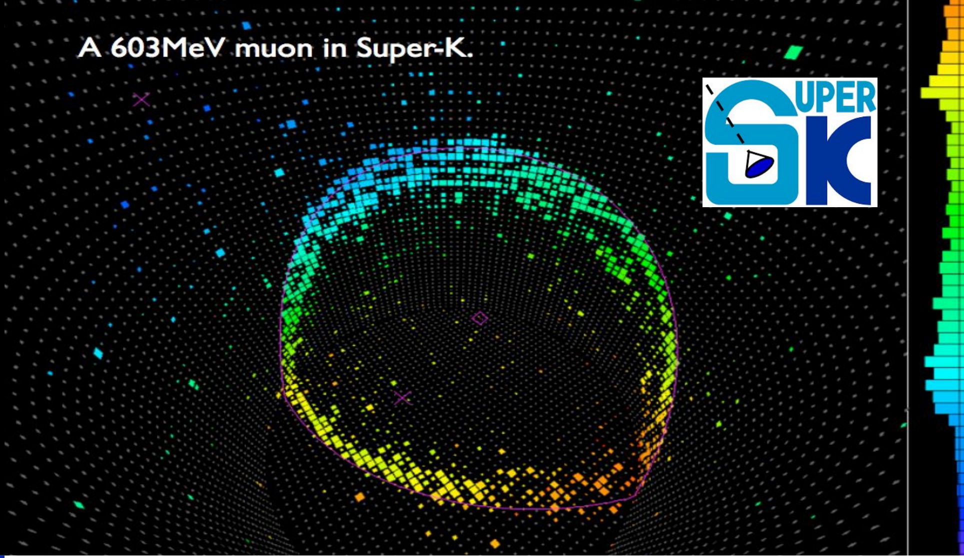
A 603MeV muon in Super-K.
SUPER K

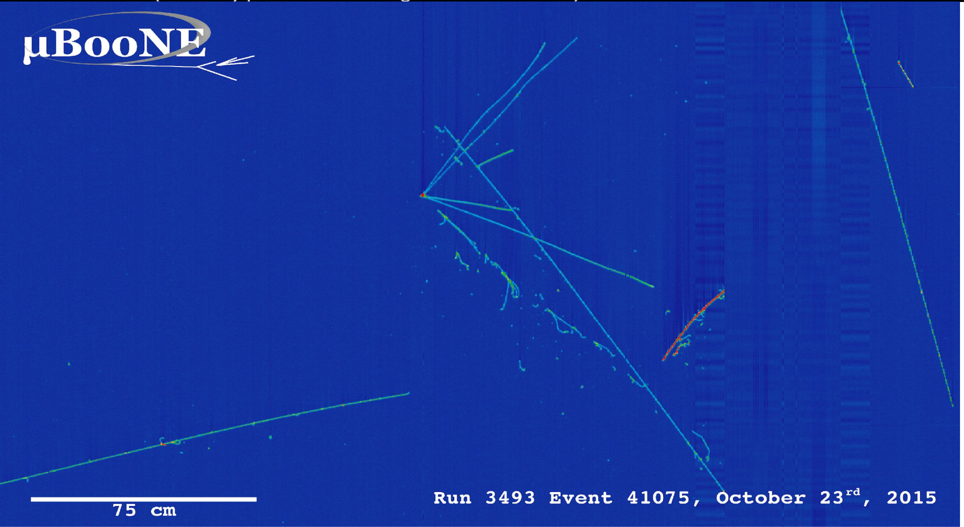
μBooNE
75 cm
Run 3493 Event 41075, October 23rd, 2015

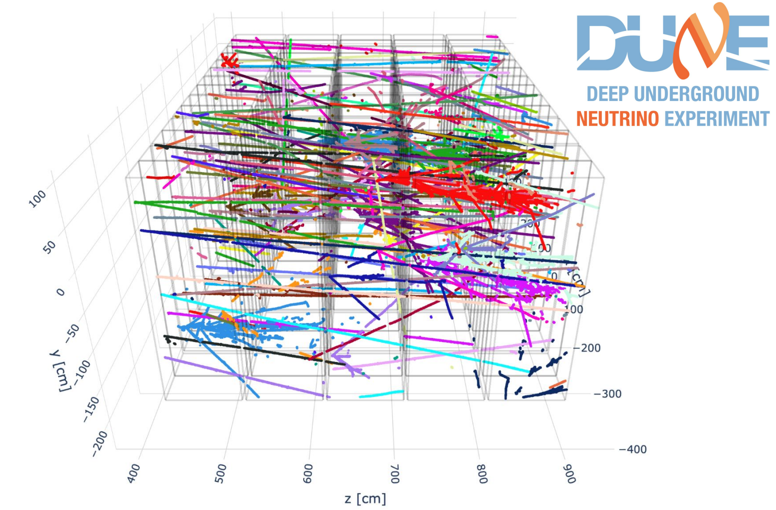
DUNE
DEEP UNDERGROUND NEUTRINO EXPERIMENT
y [cm]
z [cm]
400
500
600
700
800
900

Neutrino dataset

Rich images from large detectors
Globally sparse, locally dense

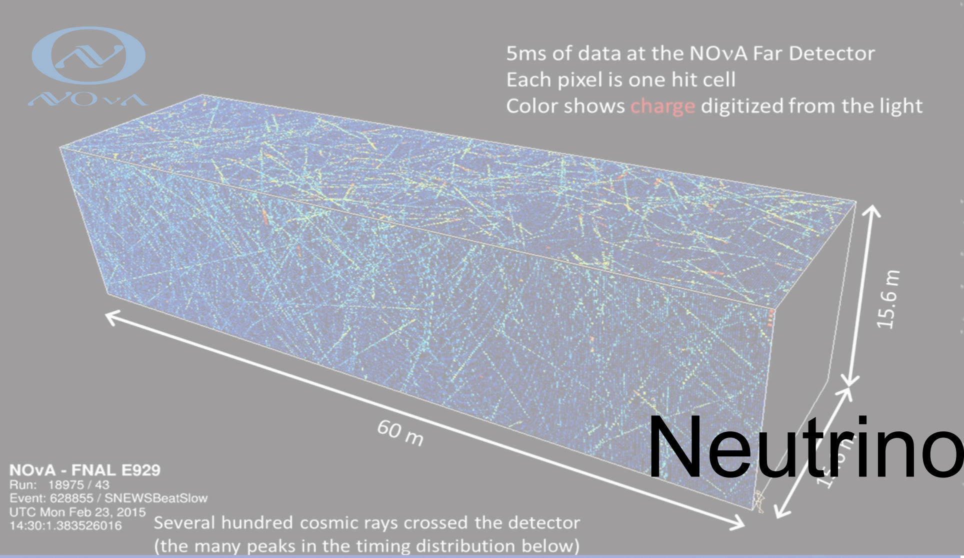

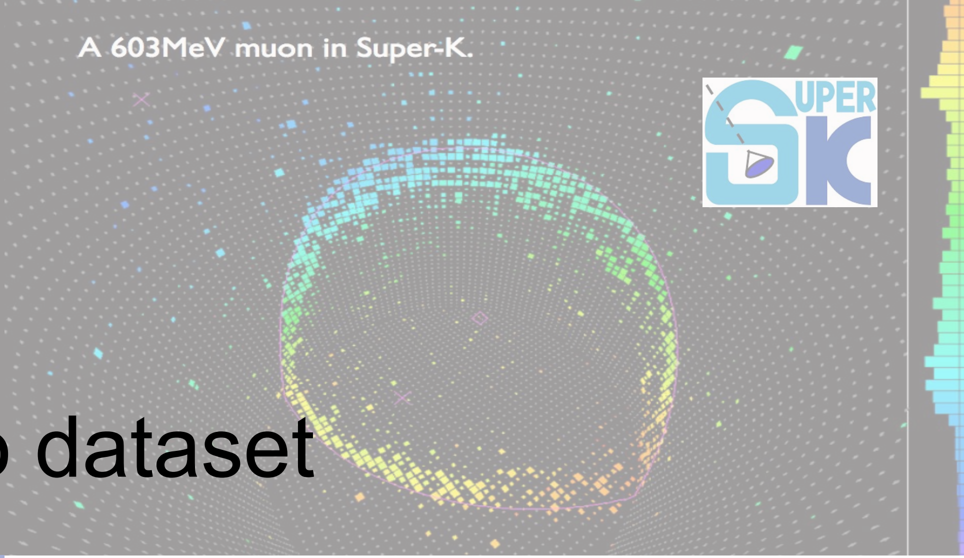

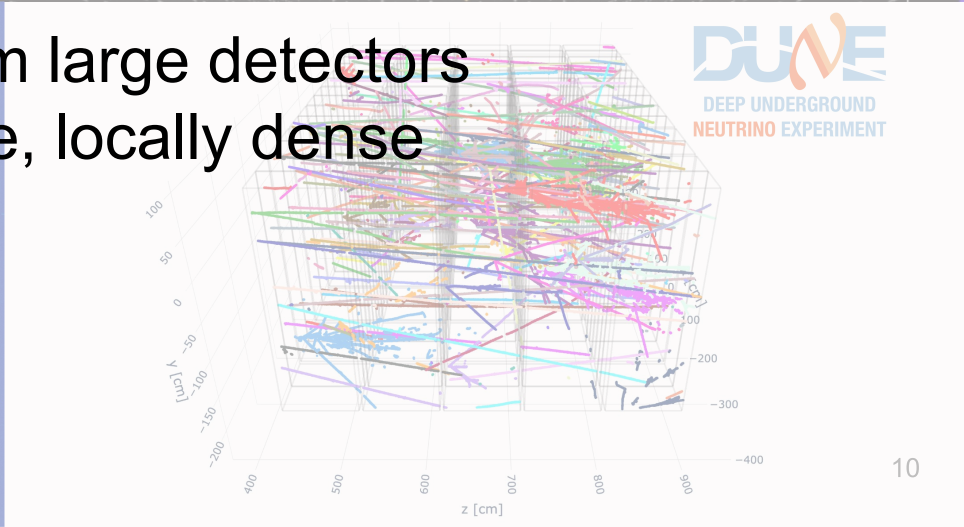

Accelerator neutrino oscillation experiments

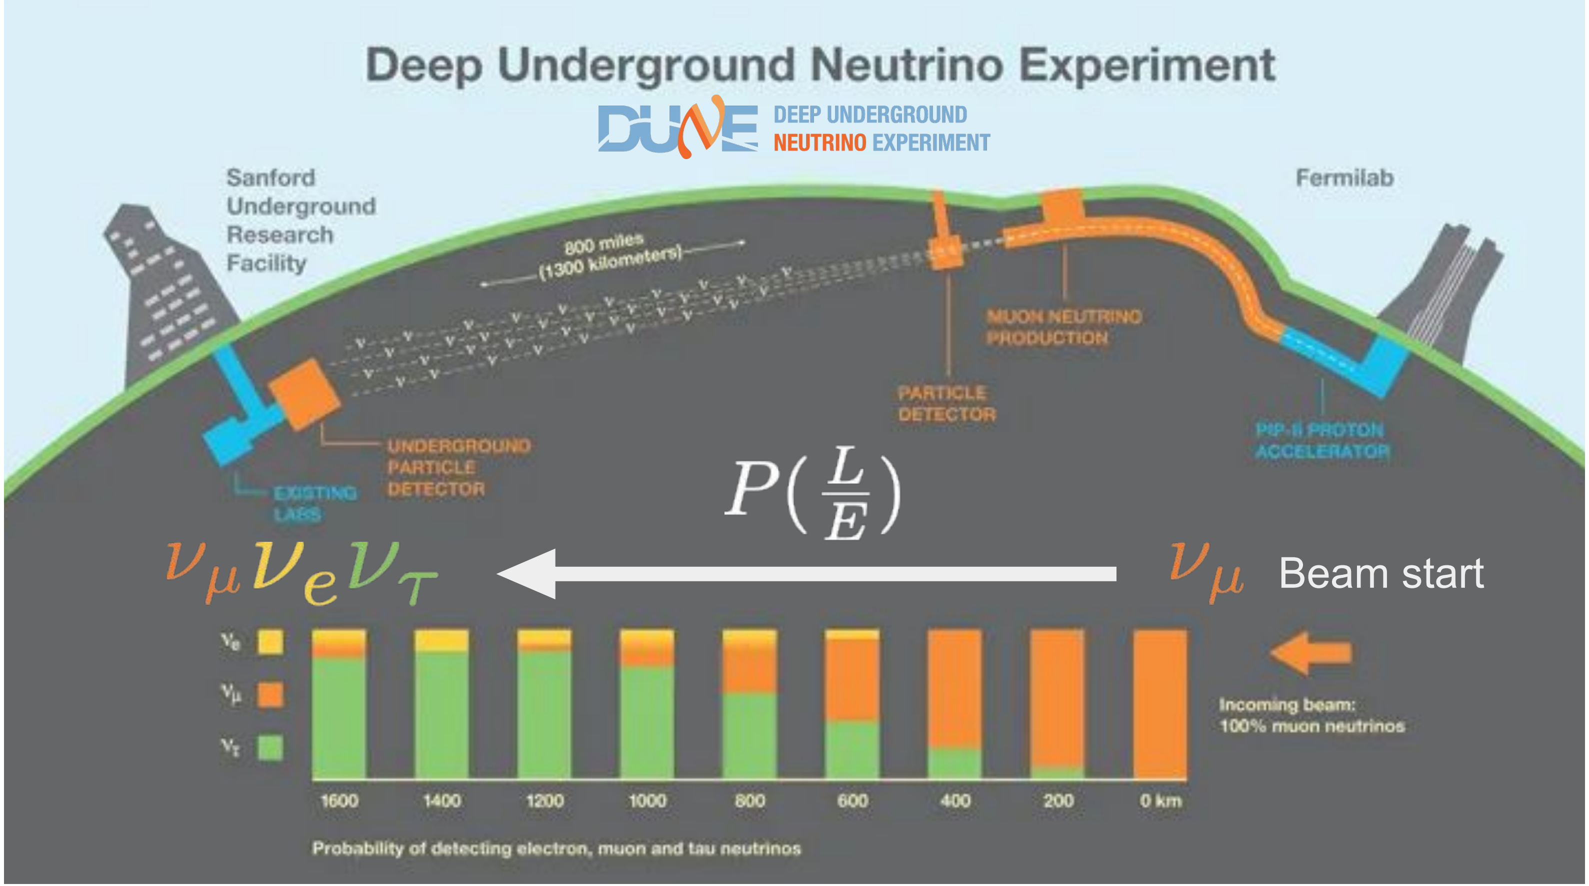

Physics analysis need

Identify neutrino flavours and energy

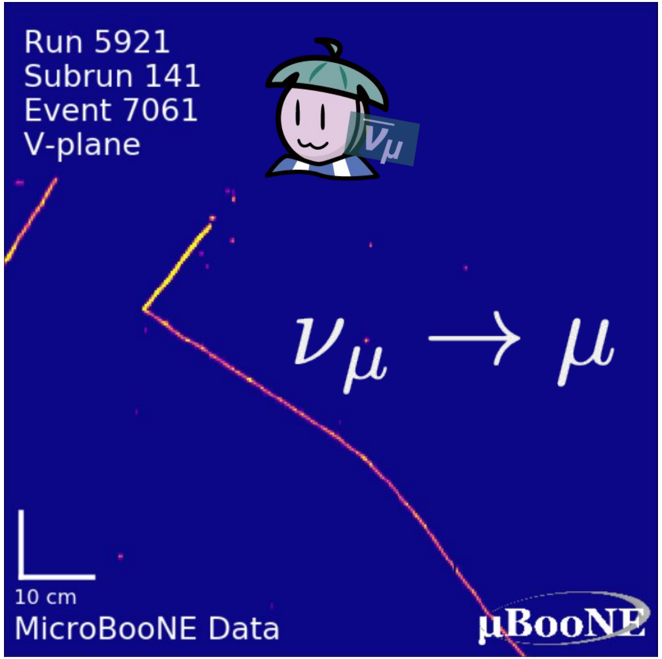

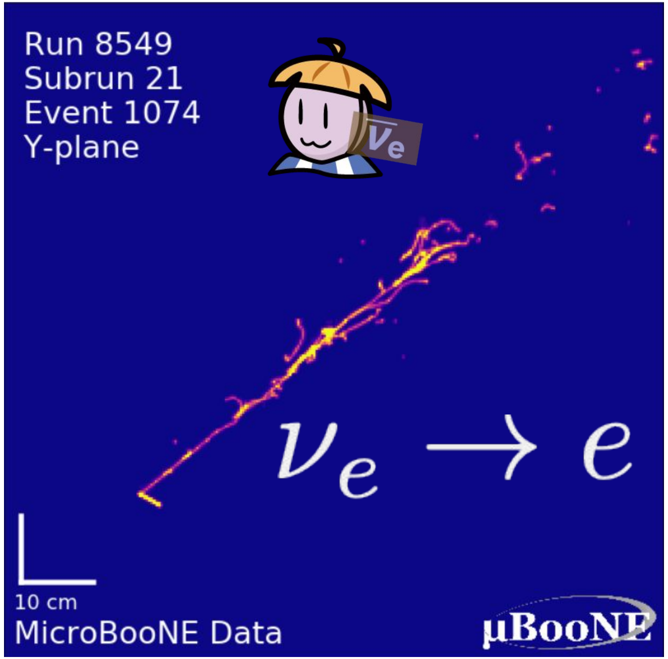

Classical
cat or dog
classification

Event reconstruction with neutrino images

CNN for neutrino flavour (classification) and energy (regression)

Phys. Rev. D 102, 092003

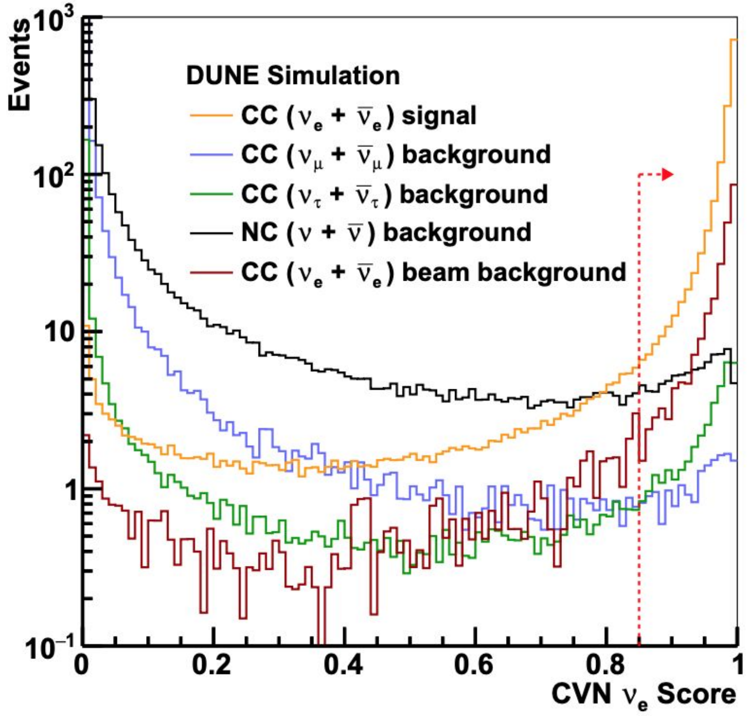

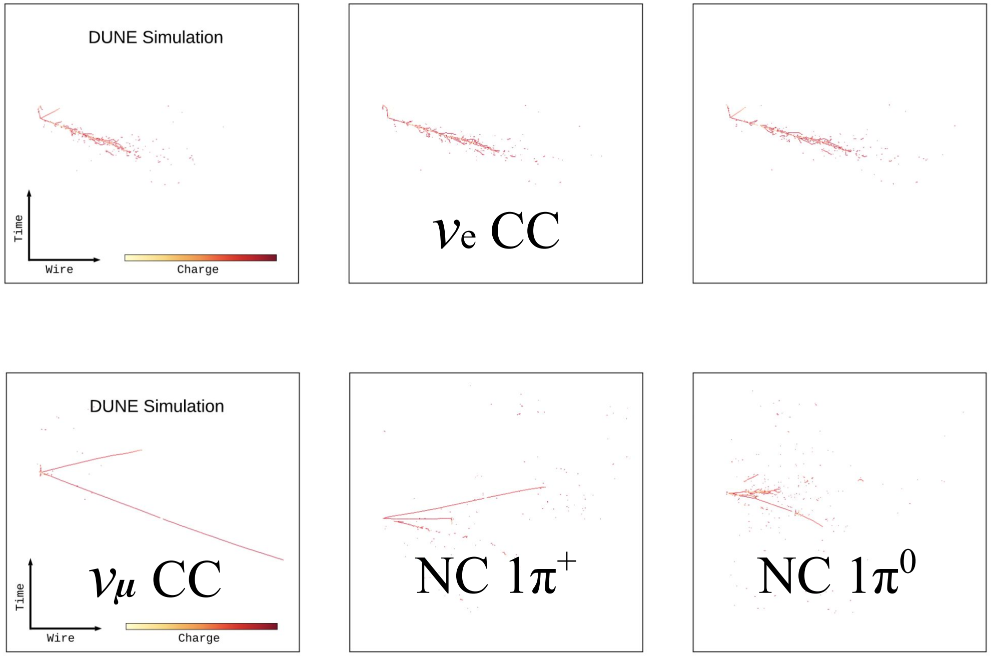

How can we improve?

Heuristic

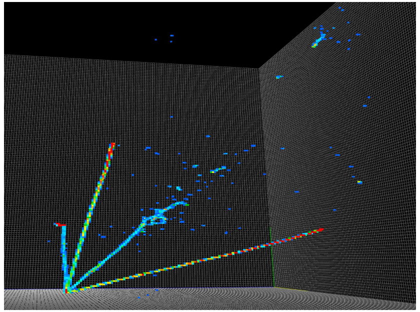

Hit formation
Interaction Clustering
Track vs Shower
Particle identification
Energy reconstruction

electron neutrino
muon neutrino

One-step

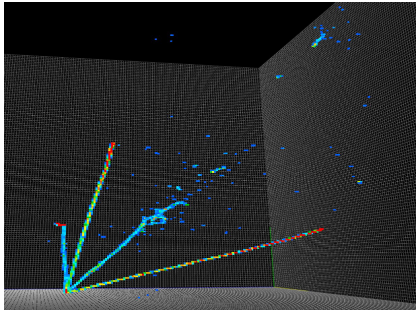

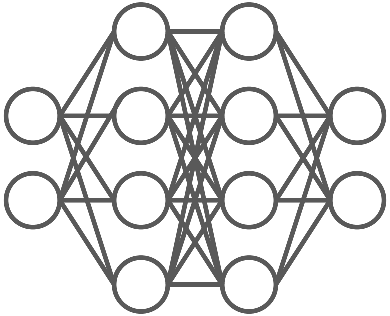

0 = electron neutrino
1 = muon neutrino

What we want: Robustness, transferability, interpretability, automation, powerful optimization methods.

End-to-end multi-step reconstruction chain

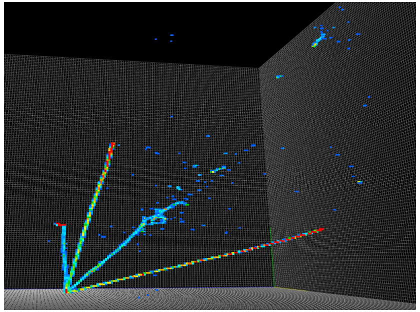

SPINE

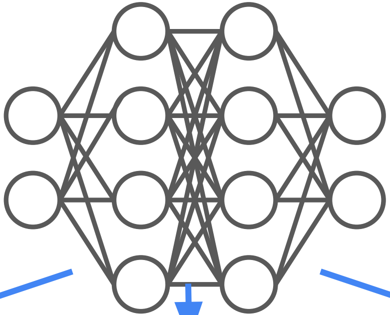

0 = electron neutrino

1 = muon neutrino

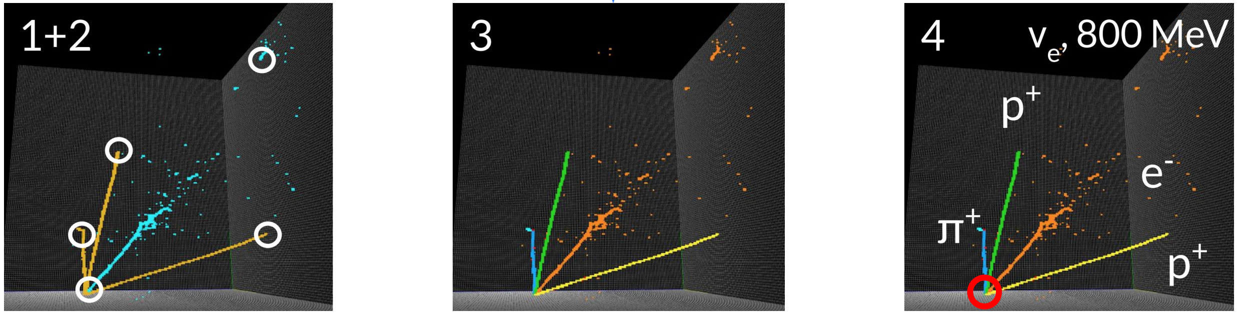

End-to-end reconstruction

with **intermediate physics observables** and **hierarchical correlation**

Scalable Particle Imaging with Neural Embedding

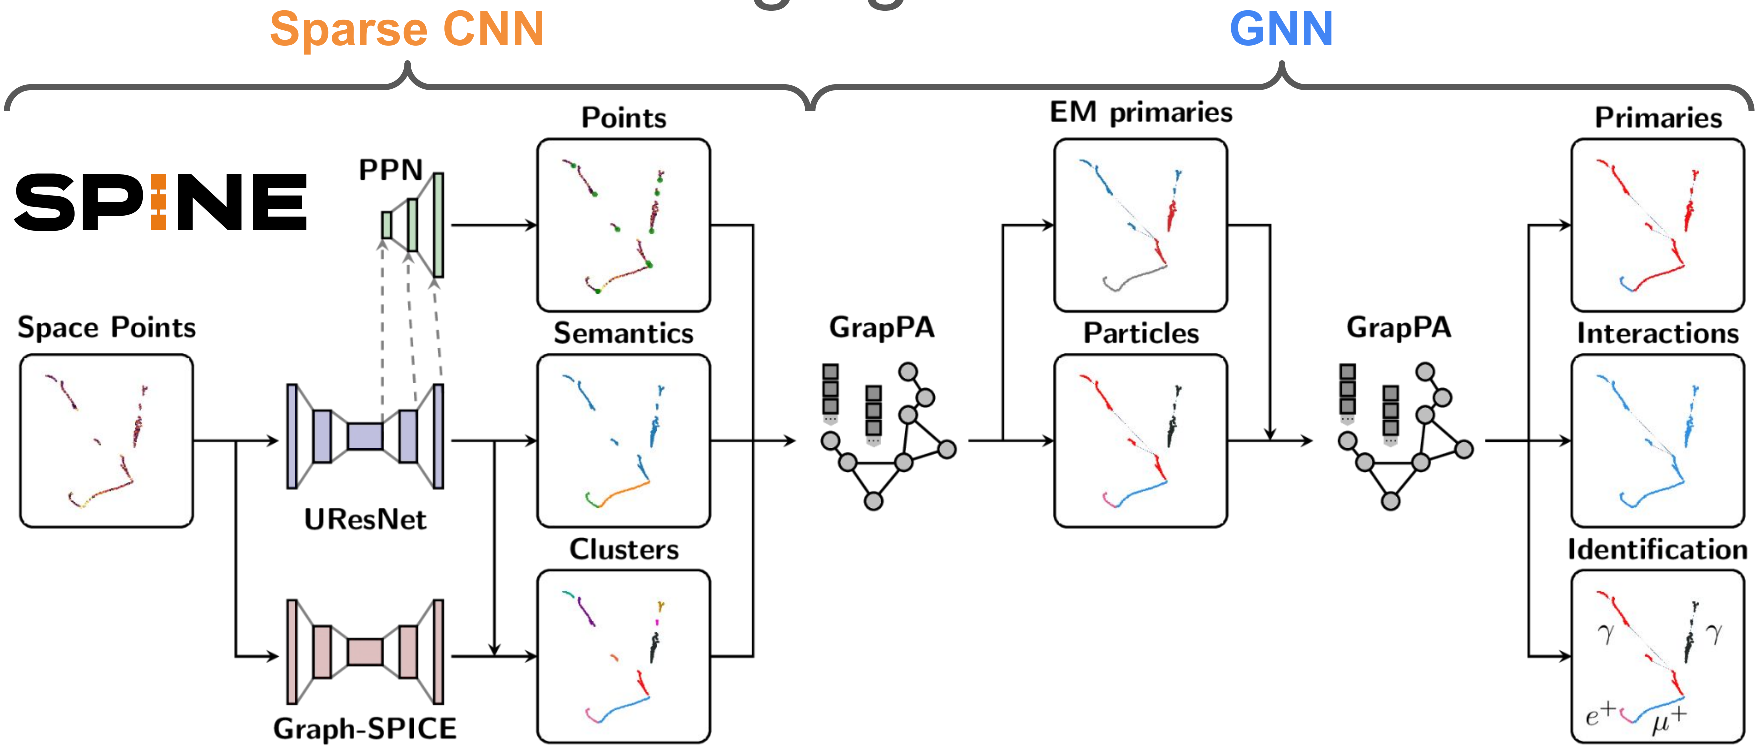

3D end-to-end optimizable data reconstruction chain:
Interpretable, optimizable, transferable, portable, HPC-compatible, accessible

Sparse CNN: Pixel-level feature extractions

UResNet (UNet + ResNet + Sparse Convolutions)

Only scale with signal pixel

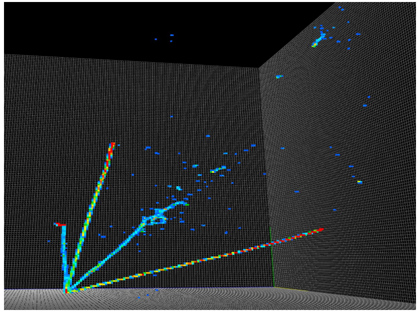

Dense

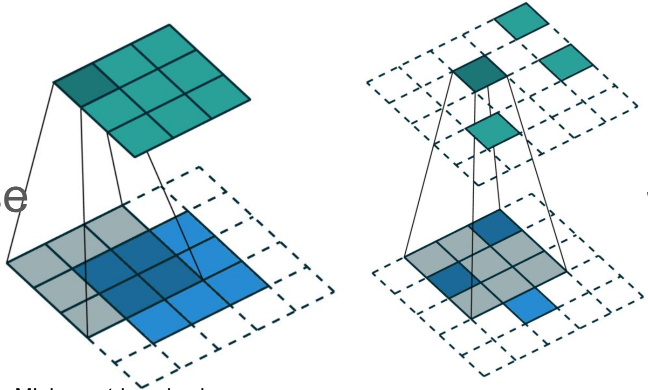

Sparse

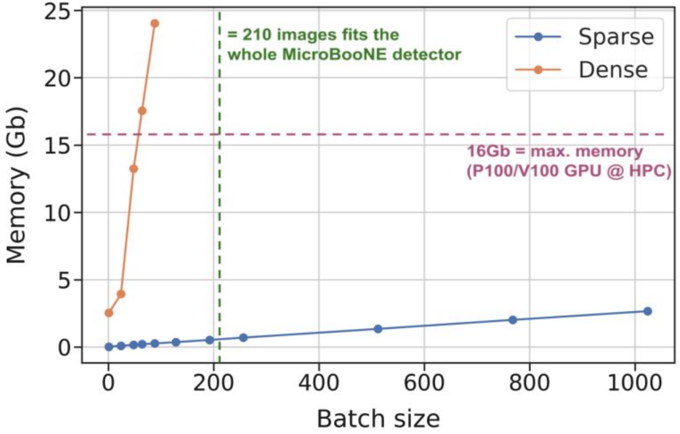

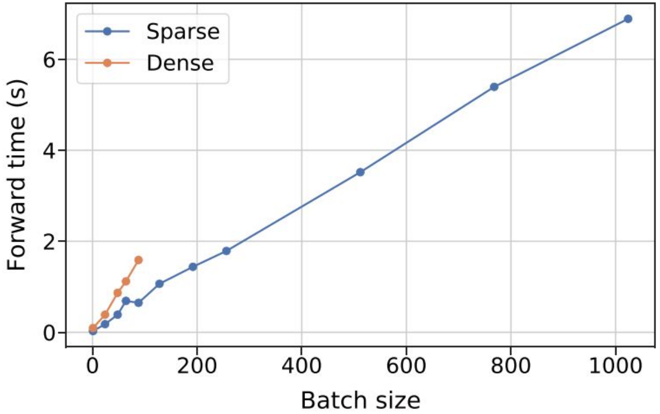

1. Semantic segmentation

Topological categories: Tracks, Showers, delta rays, Michel electrons, low energy blips

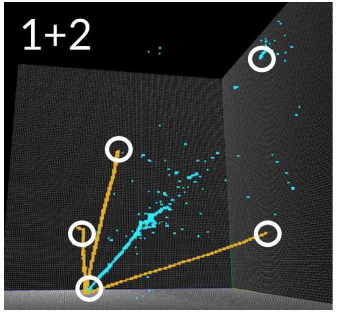

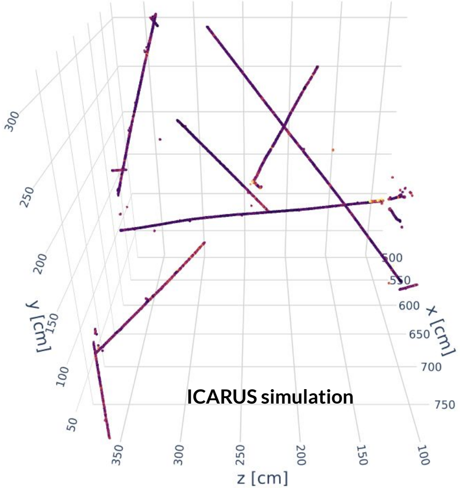

Classify pixels into categories with UResNet

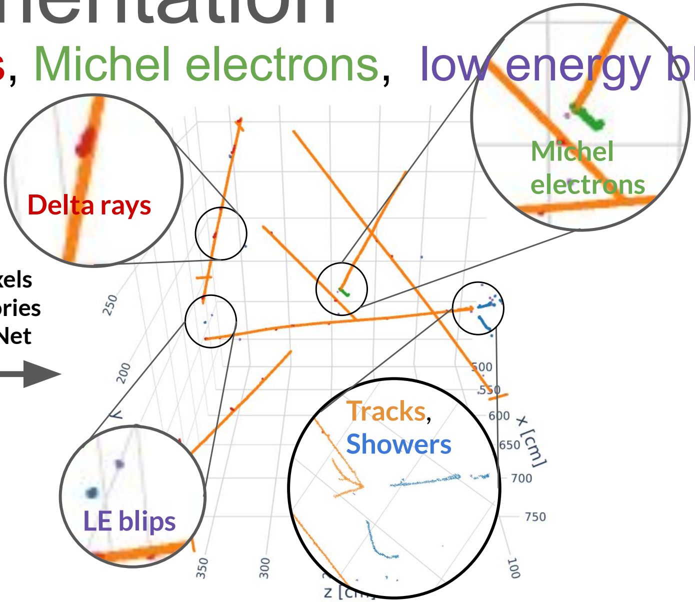

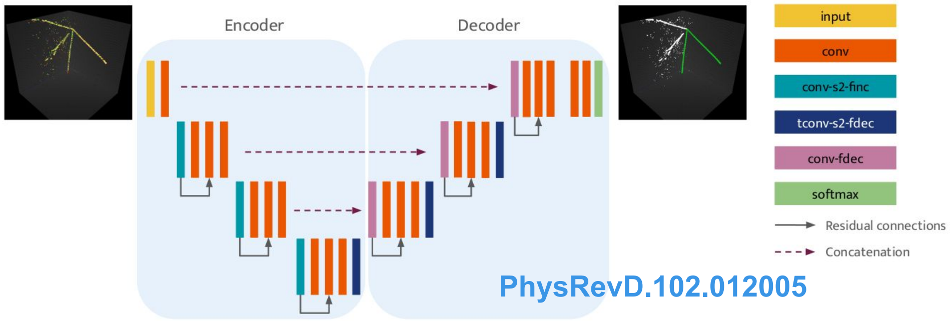

PhysRevD.102.012005

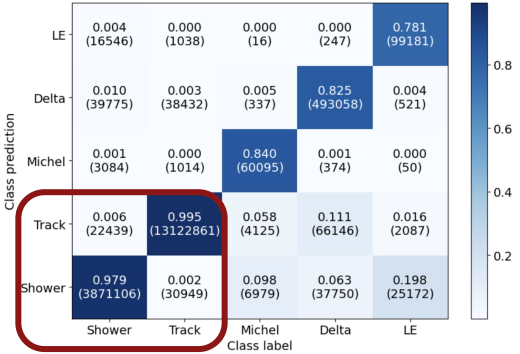

2. Point detection

Detecting points of interest (PPN):
Track starts/ends, shower fragment starts

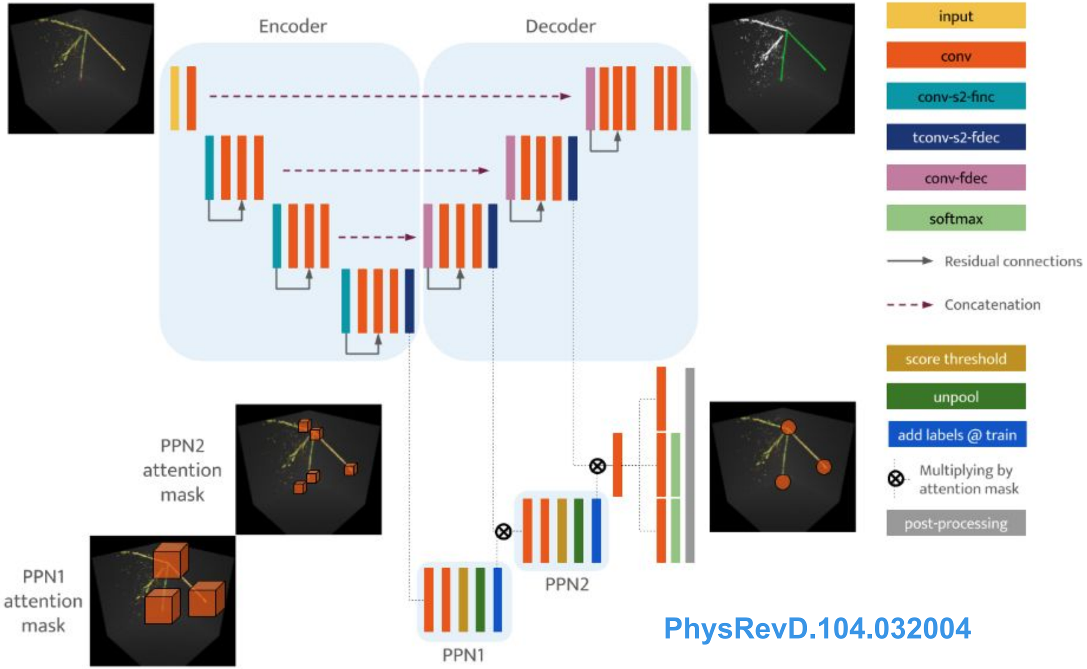

PhysRevD.104.032004

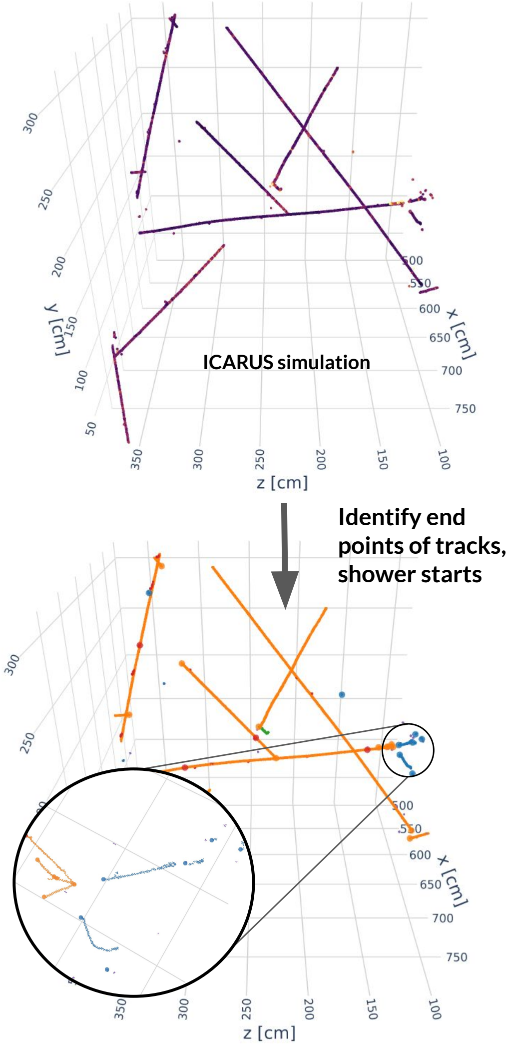

3. Clustering: Pixels to particle trajectories

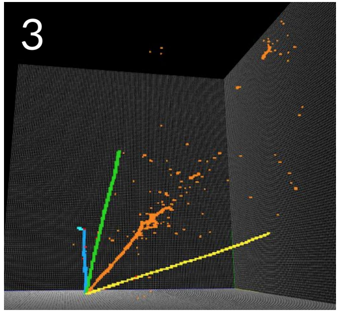

3.1: Graph-SPICE (CNN) to connect neighboring pixels into fragments
3.2: GrapPA (GNN) to aggregate fragments into particles

arXiv:2007.03083

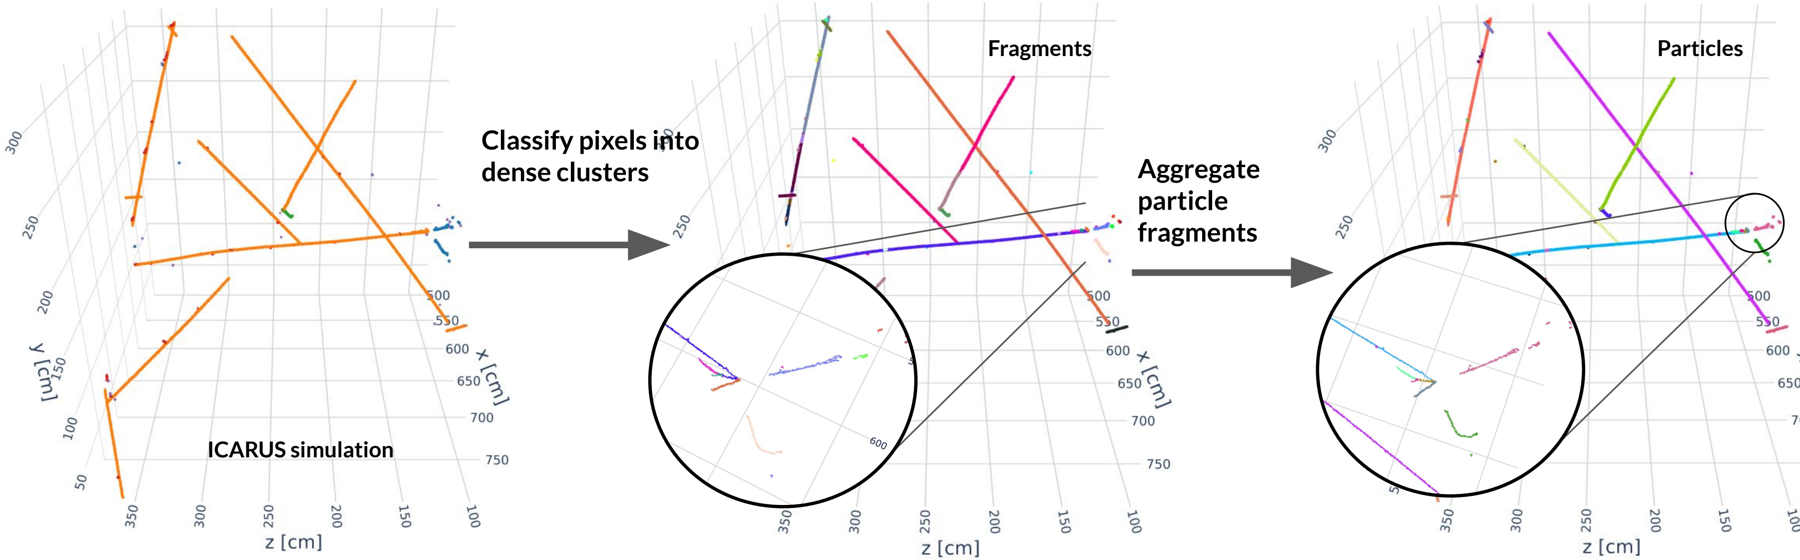

4. Identifying interactions through particle flow

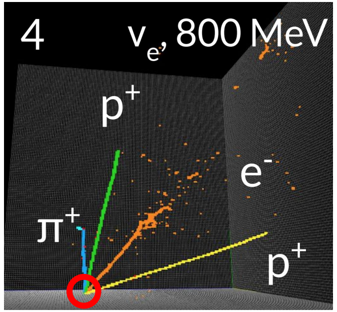

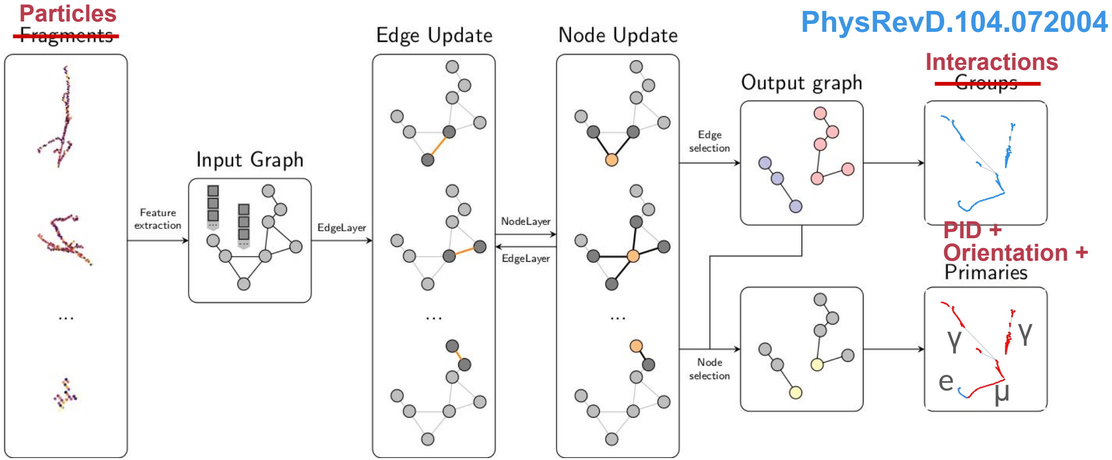

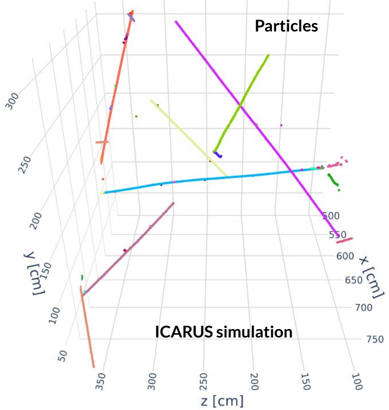

Aggregate particles

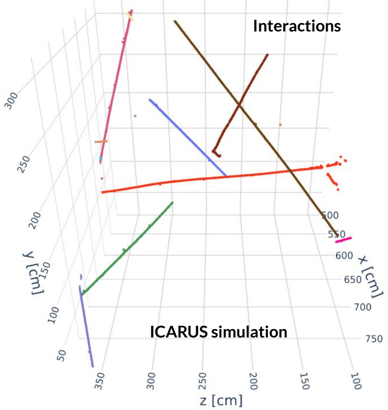

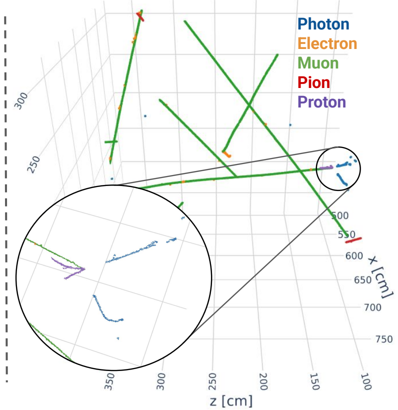

Class prediction \ Class label	Photon	Electron	Muon	Pion	Proton
Proton	0.000 (0)	0.000 (0)	0.035 (8)	0.140 (78)	0.903 (1492)
Pion	0.000 (0)	0.000 (0)	0.035 (8)	0.767 (429)	0.085 (141)
Muon	0.000 (0)	0.000 (0)	0.930 (213)	0.093 (52)	0.012 (19)
Electron	0.024 (26)	0.898 (193)	0.000 (0)	0.000 (0)	0.000 (0)
Photon	0.976 (1067)	0.102 (22)	0.000 (0)	0.000 (0)	0.000 (0)

The next-step of ML in neutrino physics

Optimize simulation from data directly

Learn representations from data directly

Supervised learnings rely on data–simulation agreement

Optimizable simulation

- Optimize on real data to improve model fidelity
- Accelerate and automate model tuning
- Use physics knowledge as much as possible
- Incorporate model correlation

Solving **inverse problems**
Detector calibration, detector design, particle reconstruction
by **Differentiable simulation** (access **gradients**)

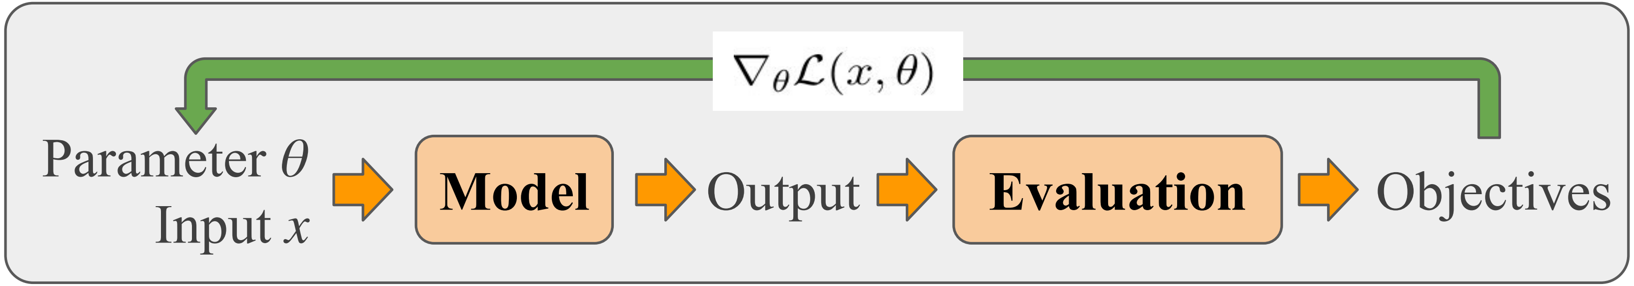

White Box: Differentiable water Cherenkov modeling

arXiv:2602.24129

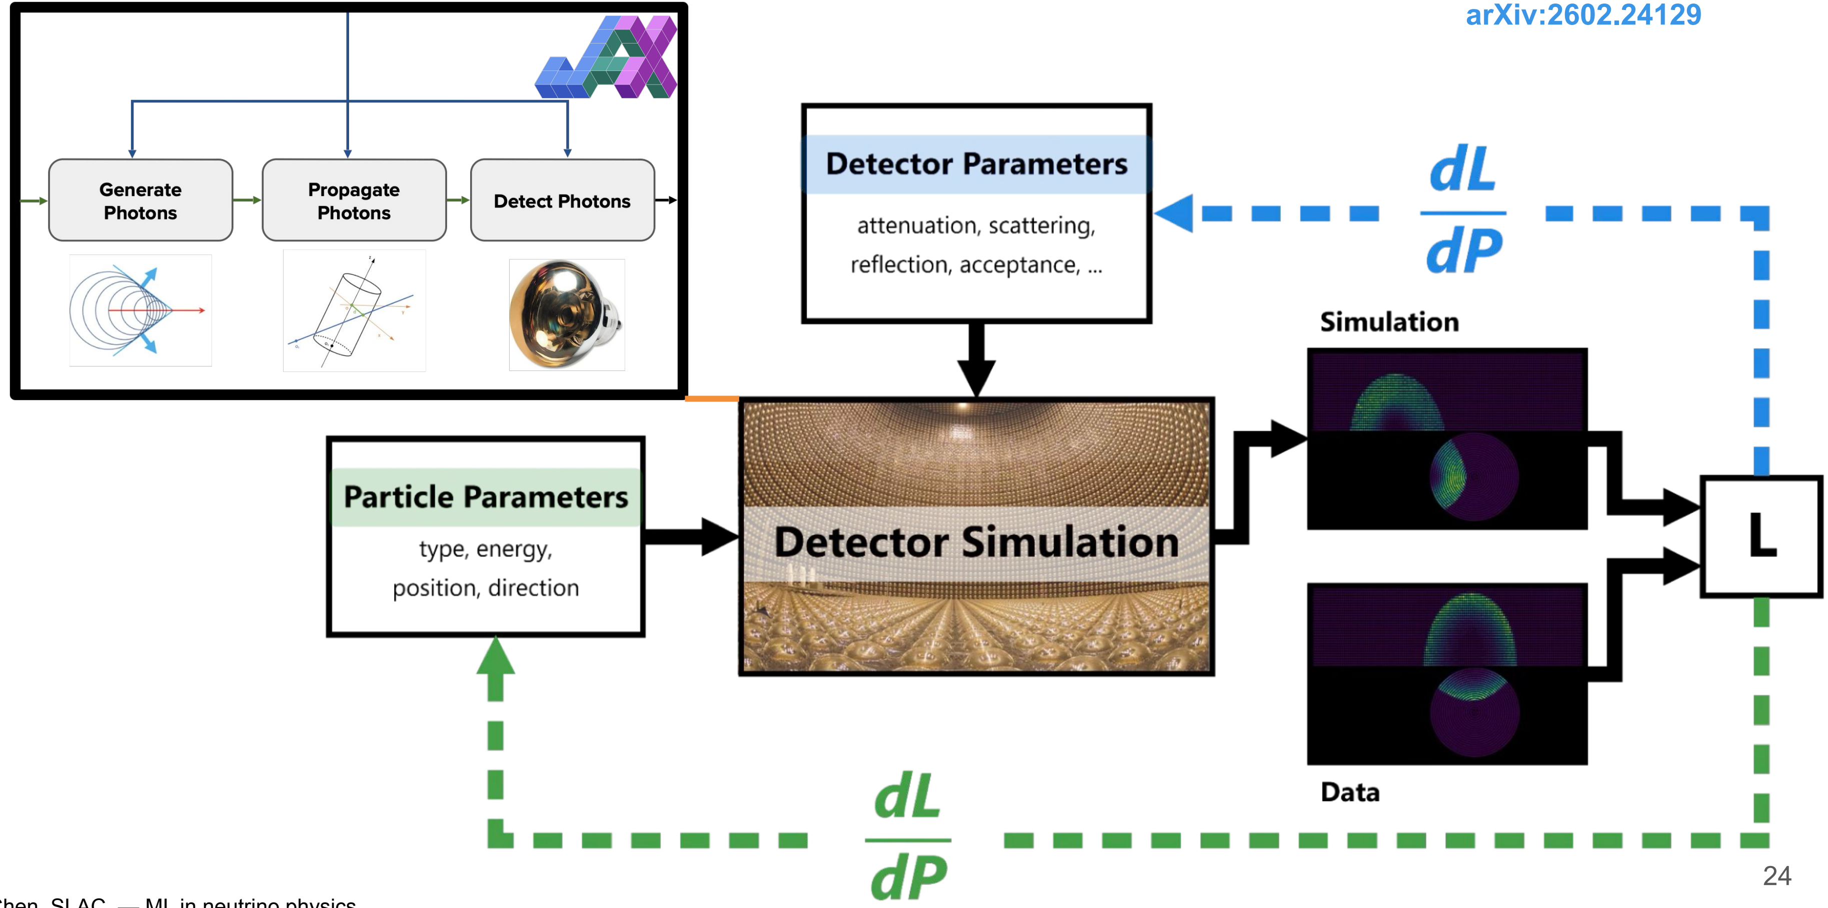

White Box: Differentiable water Cherenkov modeling

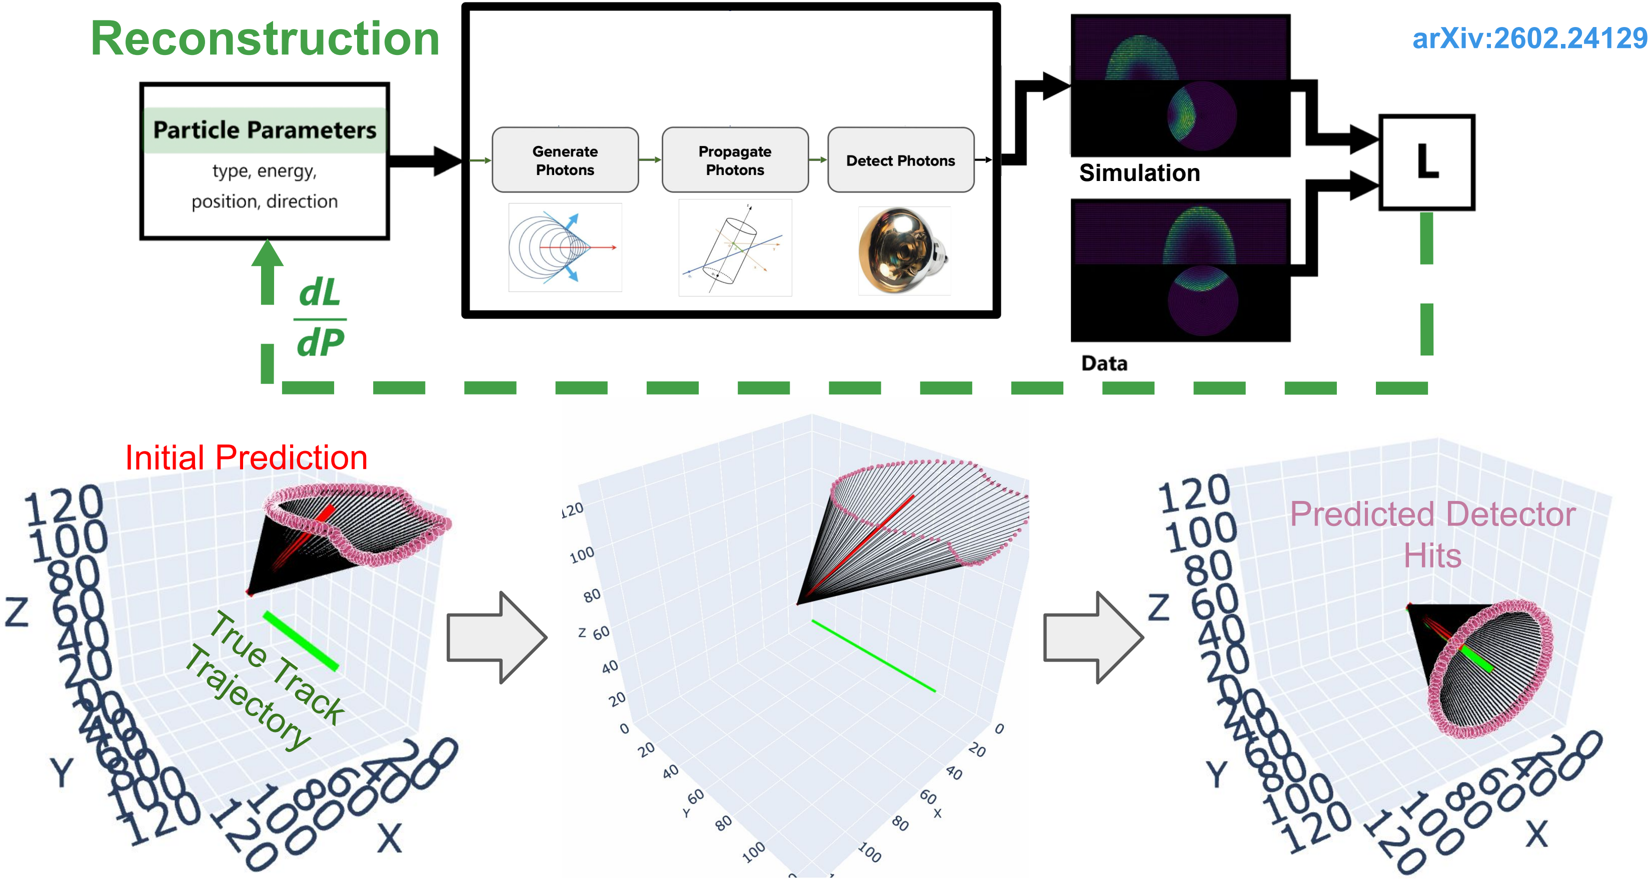

White Box: Differentiable LArTPC modeling

Mach. Learn.: Sci. Technol. 5 025012

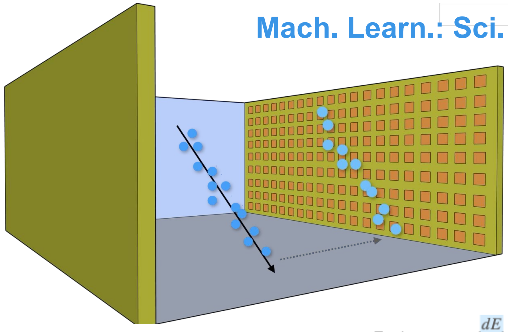

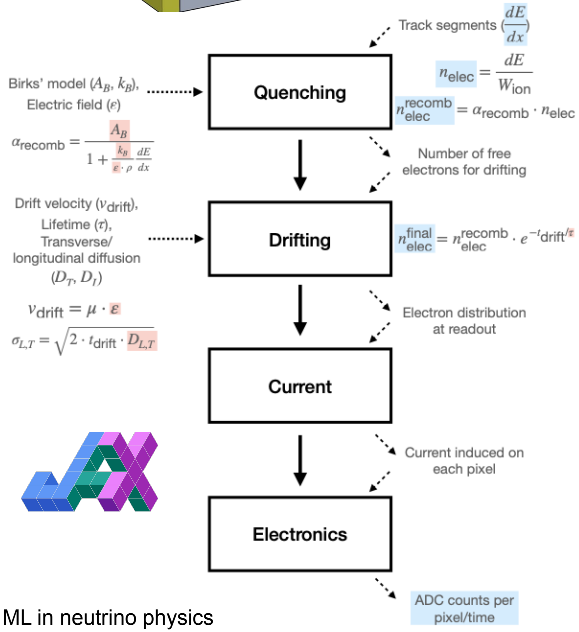

PyTorch

Calibration

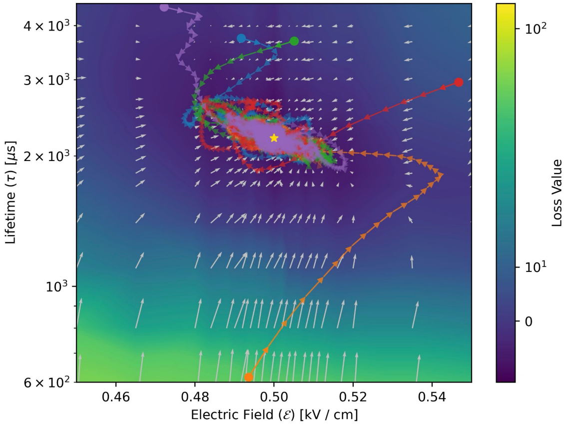

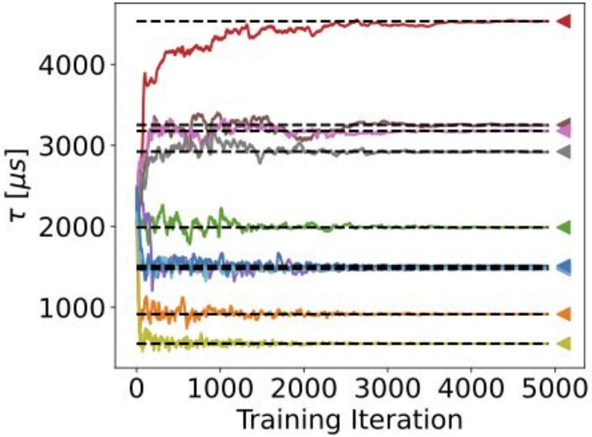

Black Box: Differentiable photon transport

- Differentiable neural surrogate for inclusive modeling
- SIREN designed to learn accurate gradient field
- Conventional photon library: discrete, sampled, not scalable
- SIREN: O(100-1000) times fewer parameters
- Optimizable directly on data

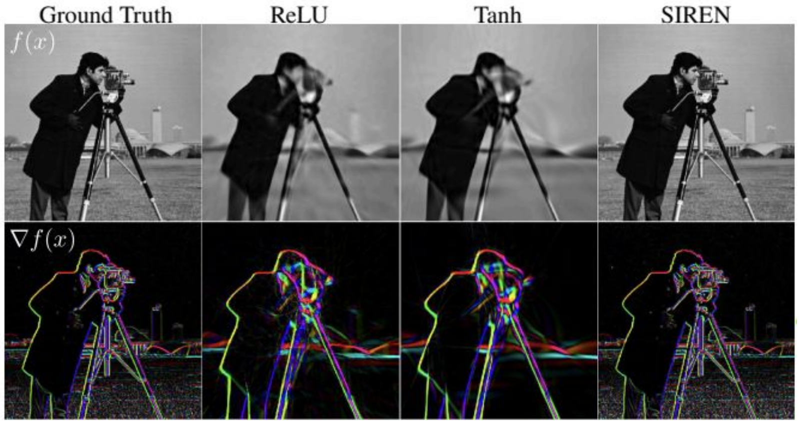

arXiv:2006.09661

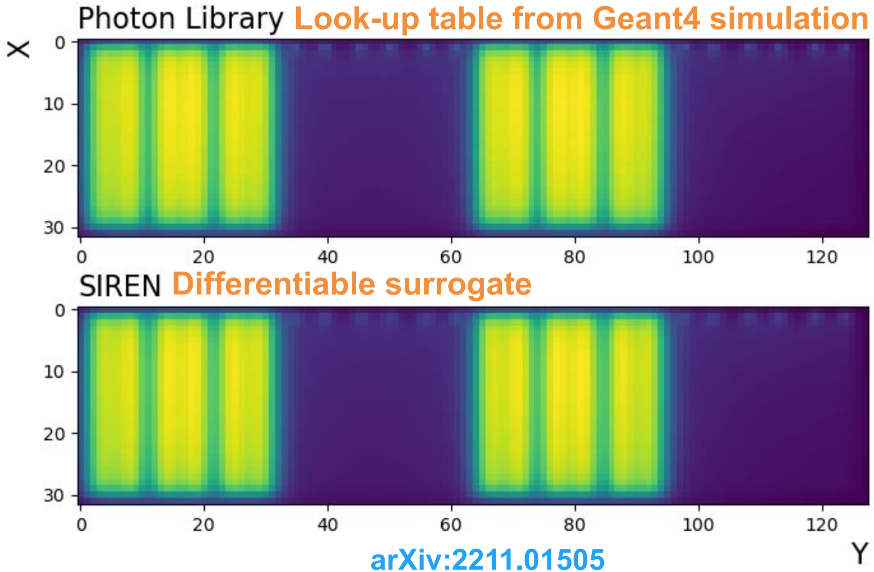

arXiv:2211.01505

Generative models

- Fast simulation for inference of LArTPC detector readout images
- Score-based diffusion model
- Indistinguishable by CNN discriminator, retain physics distribution
- Conditional (on proton momentum) latent diffusion model
- An automated inference loop!

arXiv:2307.13687

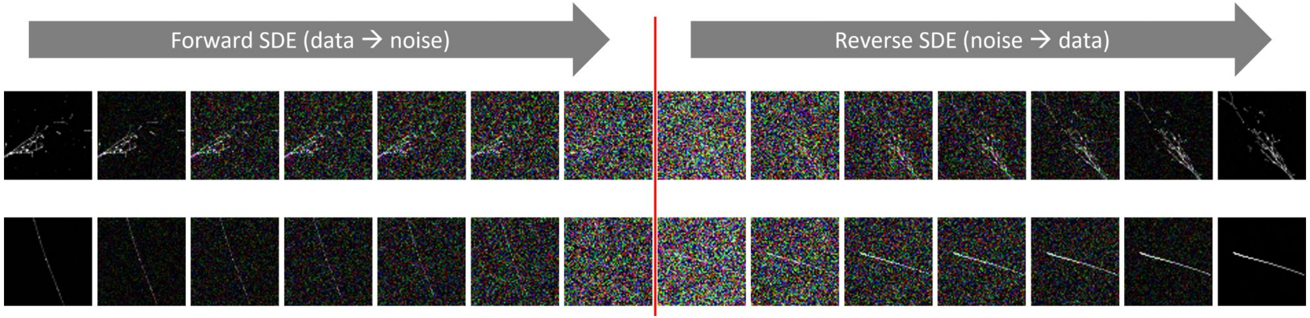

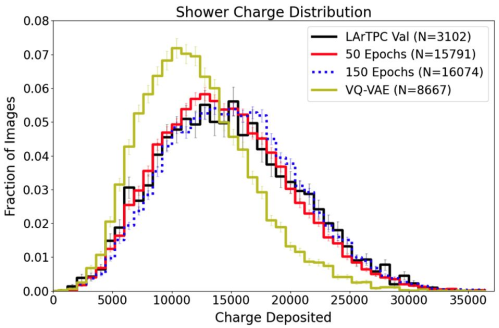

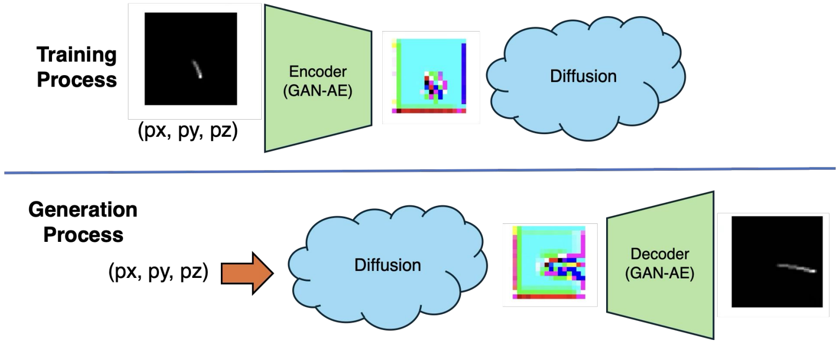

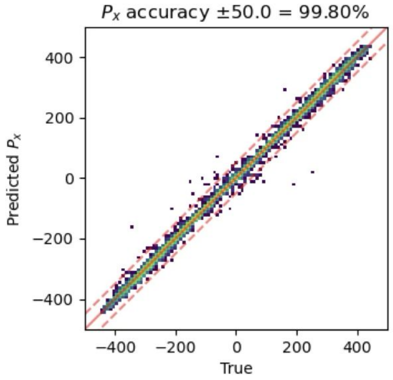

Zev Imani
NPML 2025

The next-step of ML in neutrino physics

Optimize simulation from data directly

Learn representations from data directly

Supervised learnings rely on data–simulation agreement

Self-supervised learning (SSL)

- Pre-training (directly on data) + fine-tuning (can be directly on data)
- Learn general representations from the data world = reusable features

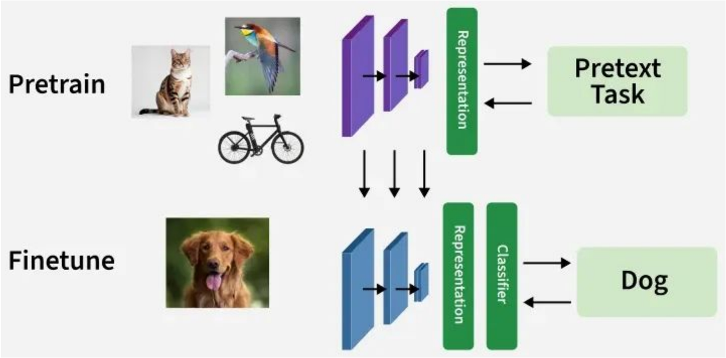

Symbolic data
Track/shower, particles
A data unit carries strong semantic meaning

Sensor data
3D point, waveform tick
A data unit carries little information

Contrastive learning

Use physics inspired **augmentation** to make representations robust against **data–simulation mismatches** and **detector uncertainties**.

arXiv:2502.07724

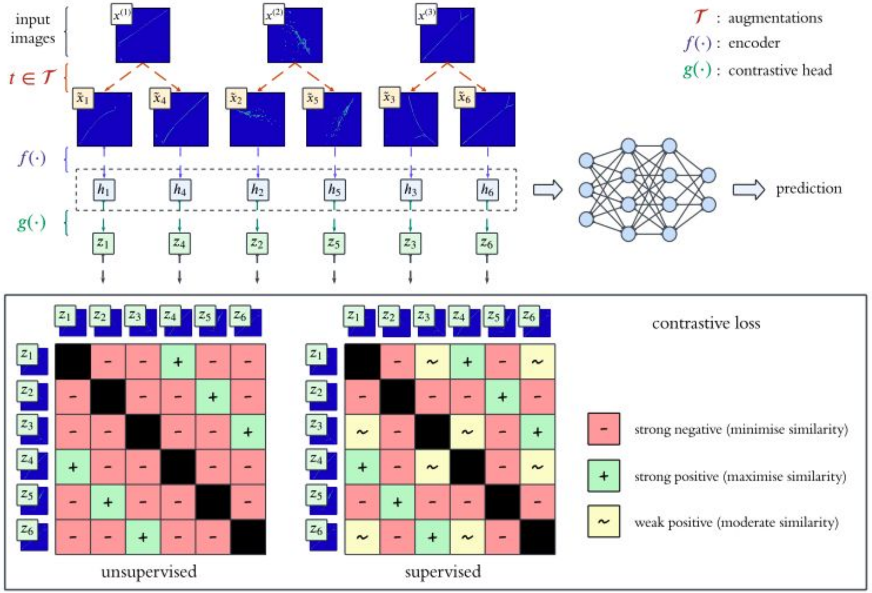

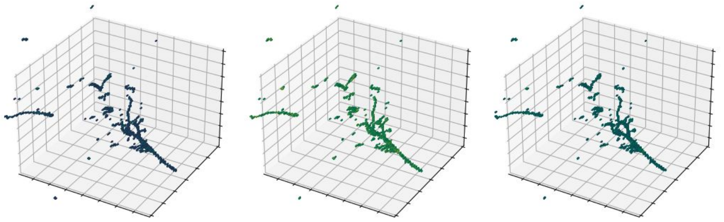

(a) Nominal (b) Throw 1 (c) Throw 2

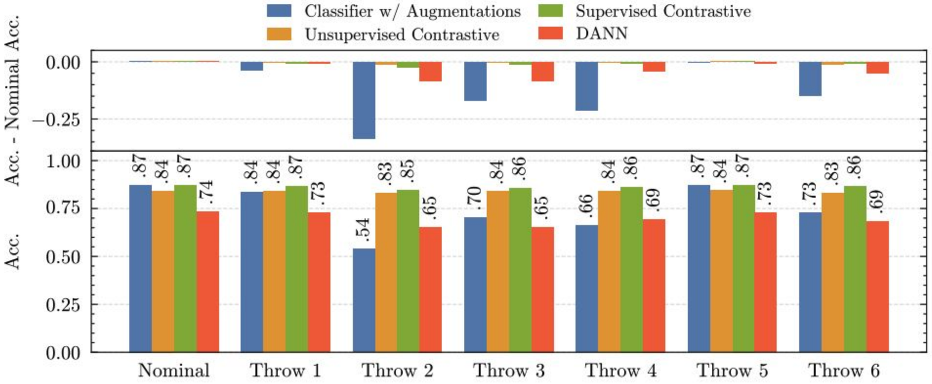

Contrastive clustering

Inspired by LZ's anomaly detection work (LZ: dark matter experiment)

Unsupervised contrastive clustering to discover unmodeled events

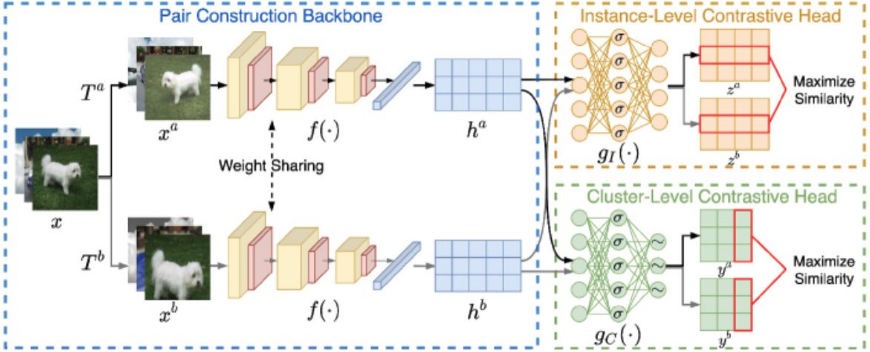

arXiv:2009.09687

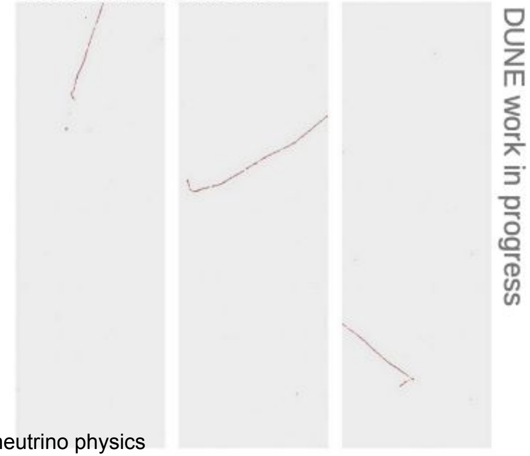

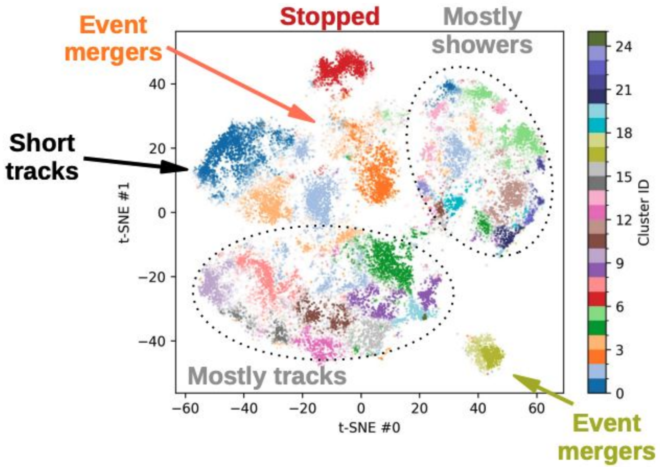

Callum Wilkinson
NPML 2025

PoLAr-MAE: Point-based LAr Masked Autoencoder

Mach. Learn.: Sci. Technol. 7 025023

Toward foundation models in HEP:

- Mask-based SSL on 3D point cloud
- Inspected learnt representations for particle reconstruction

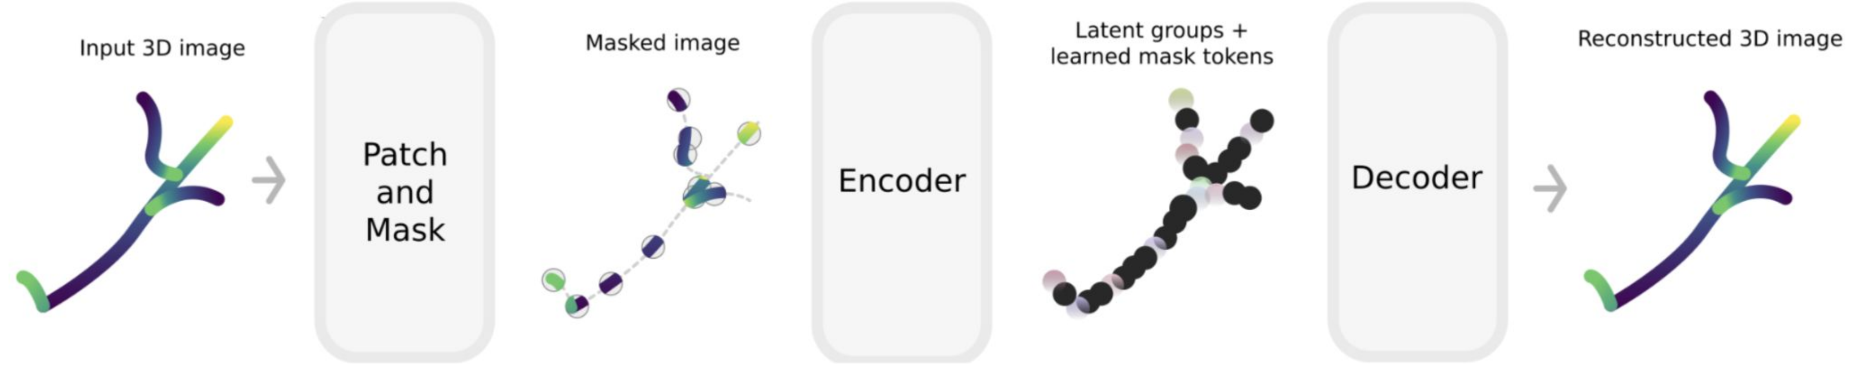

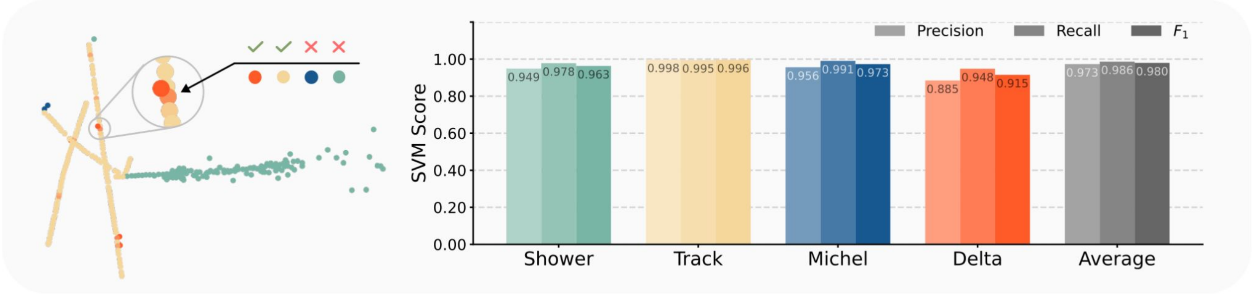

Instance Sharing

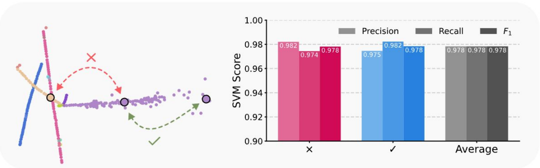

Vertex Sharing

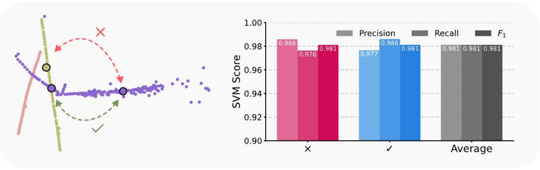

Panda: Self-distillation of reusable sensor-level representations

Toward foundation models in HEP: Test on particle reconstruction

arXiv:2512.01324

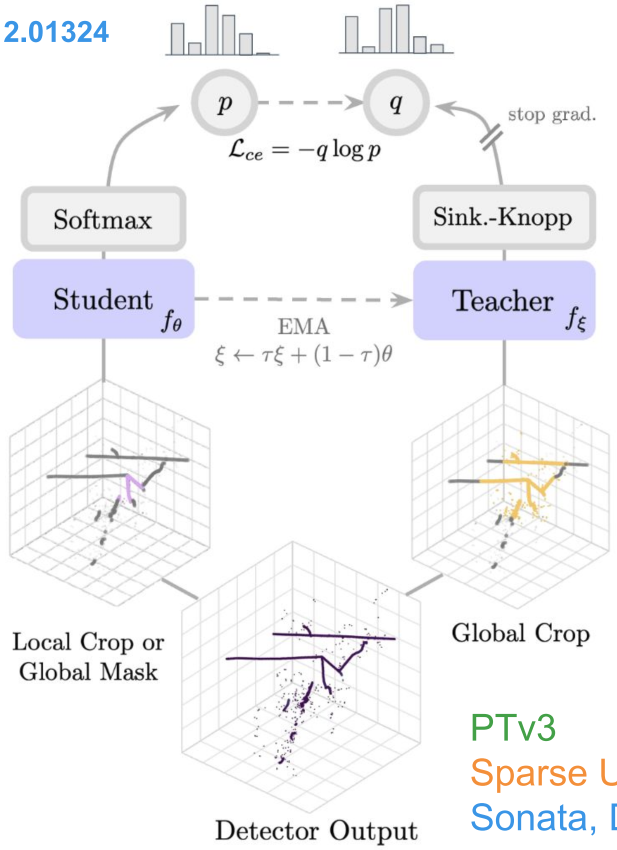

PTv3
Sparse UResNet
Sonata, DINOv2

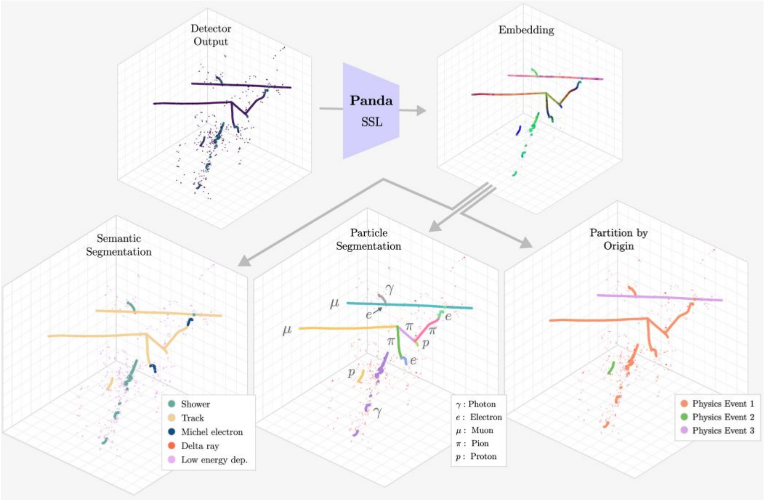

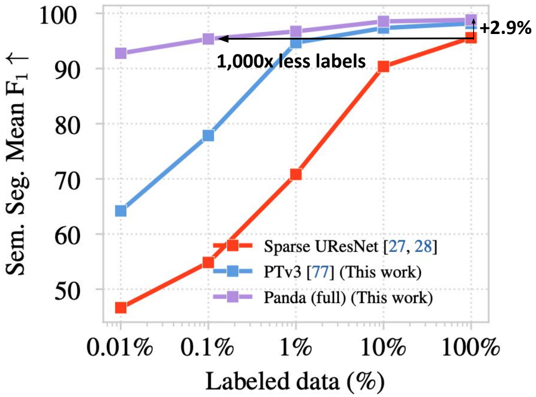

Fine tune with
fewer labels/
higher performance

From collider jets to neutrino interactions

OmniLearned: a foundation model with PT v2 as the backbone

Tasks on MINERvA: 1. Energy regression 2. Charged-current pion classification

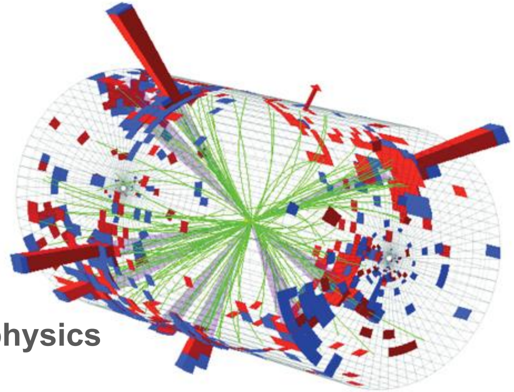

open data
TeV jets w/ QCD physics

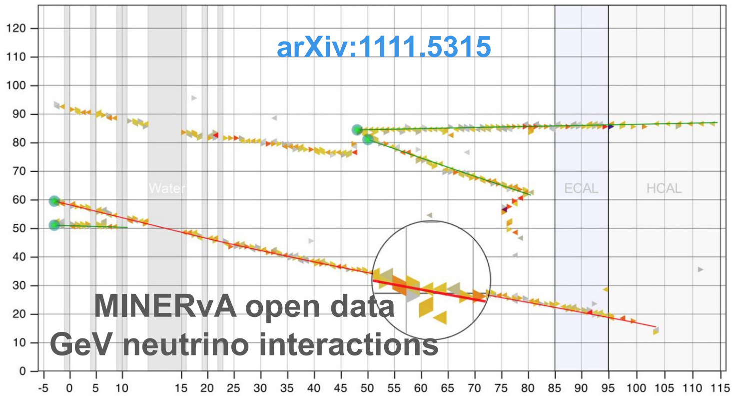

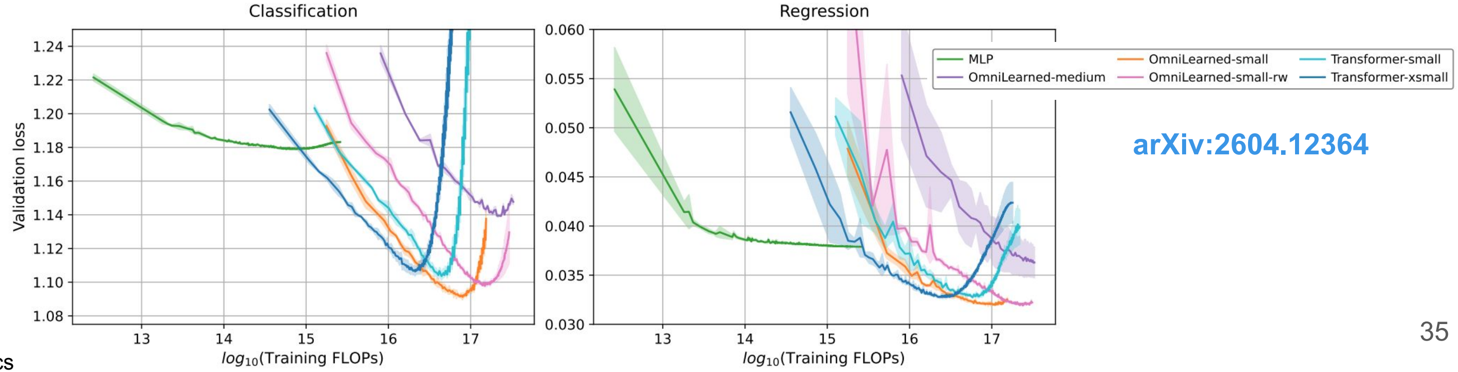

arXiv:2604.12364

Building a community ecosystem

Neutrinos Open Data
A portal to open data from neutrino experiments

NUISANCE

Workshop — Idea generation

School — Training

Open data challenge — Open data curation, Baseline solution

Knowledge hub — Leaderboard

Conference — Awards

Workshops and schools

Cross-domain AI/ML research in HEP, unconference style, across career stages

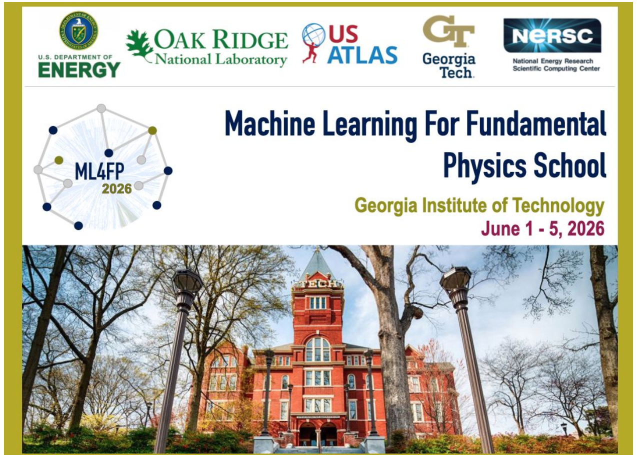

Machine Learning for Fundamental Physics School (ML4FP) 2026

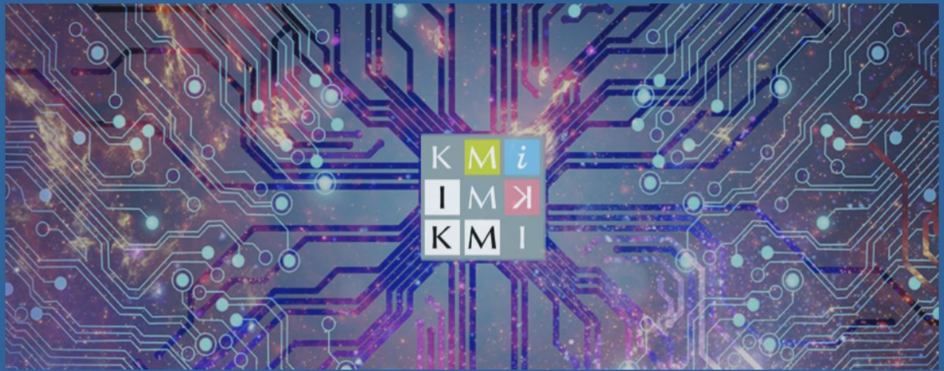

Future of Artificial Intelligence for Science in Japan (FAIRS Japan 2024)




Neutrino Physics and Machine Learning (NPML)

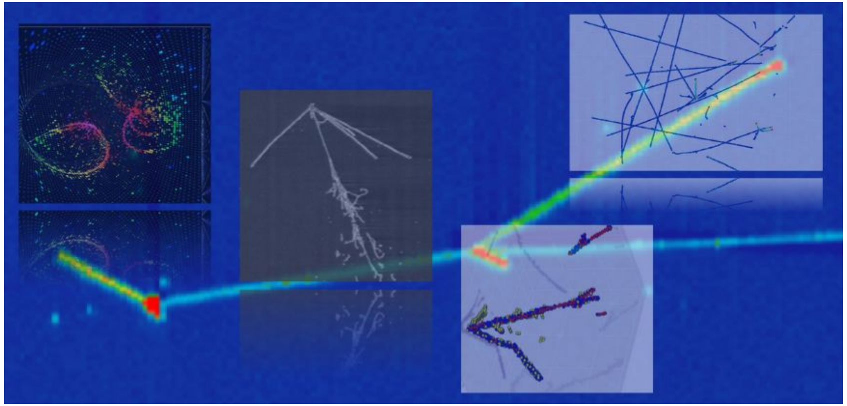

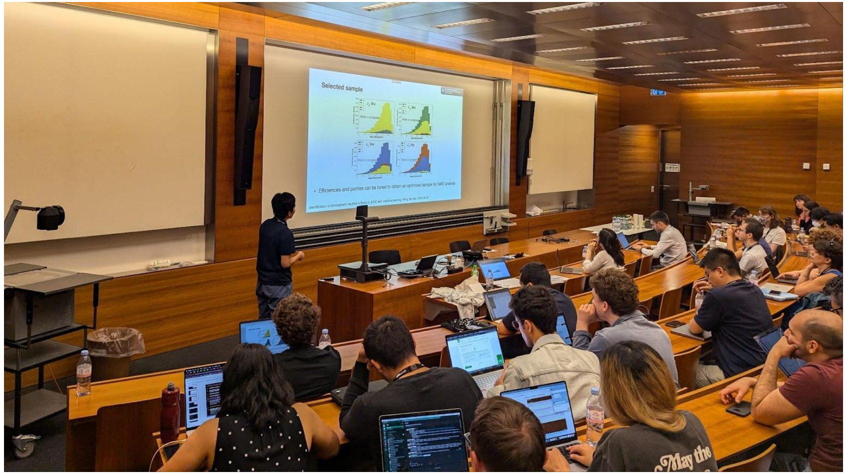

- NPML 2026, UC Irvine ←
- NPML 2025, Utokyo
- NPML 2024, ETH Zurich
- NPML 2023, Tufts
- NPML 2020, Fermilab virtual

June 15-19!
Right before NEUTRINO

More at the neutrino parallel session

Join the neutrino discussion!

Wednesday June 3rd, 15:30 - 17:00

Tell us what you do!

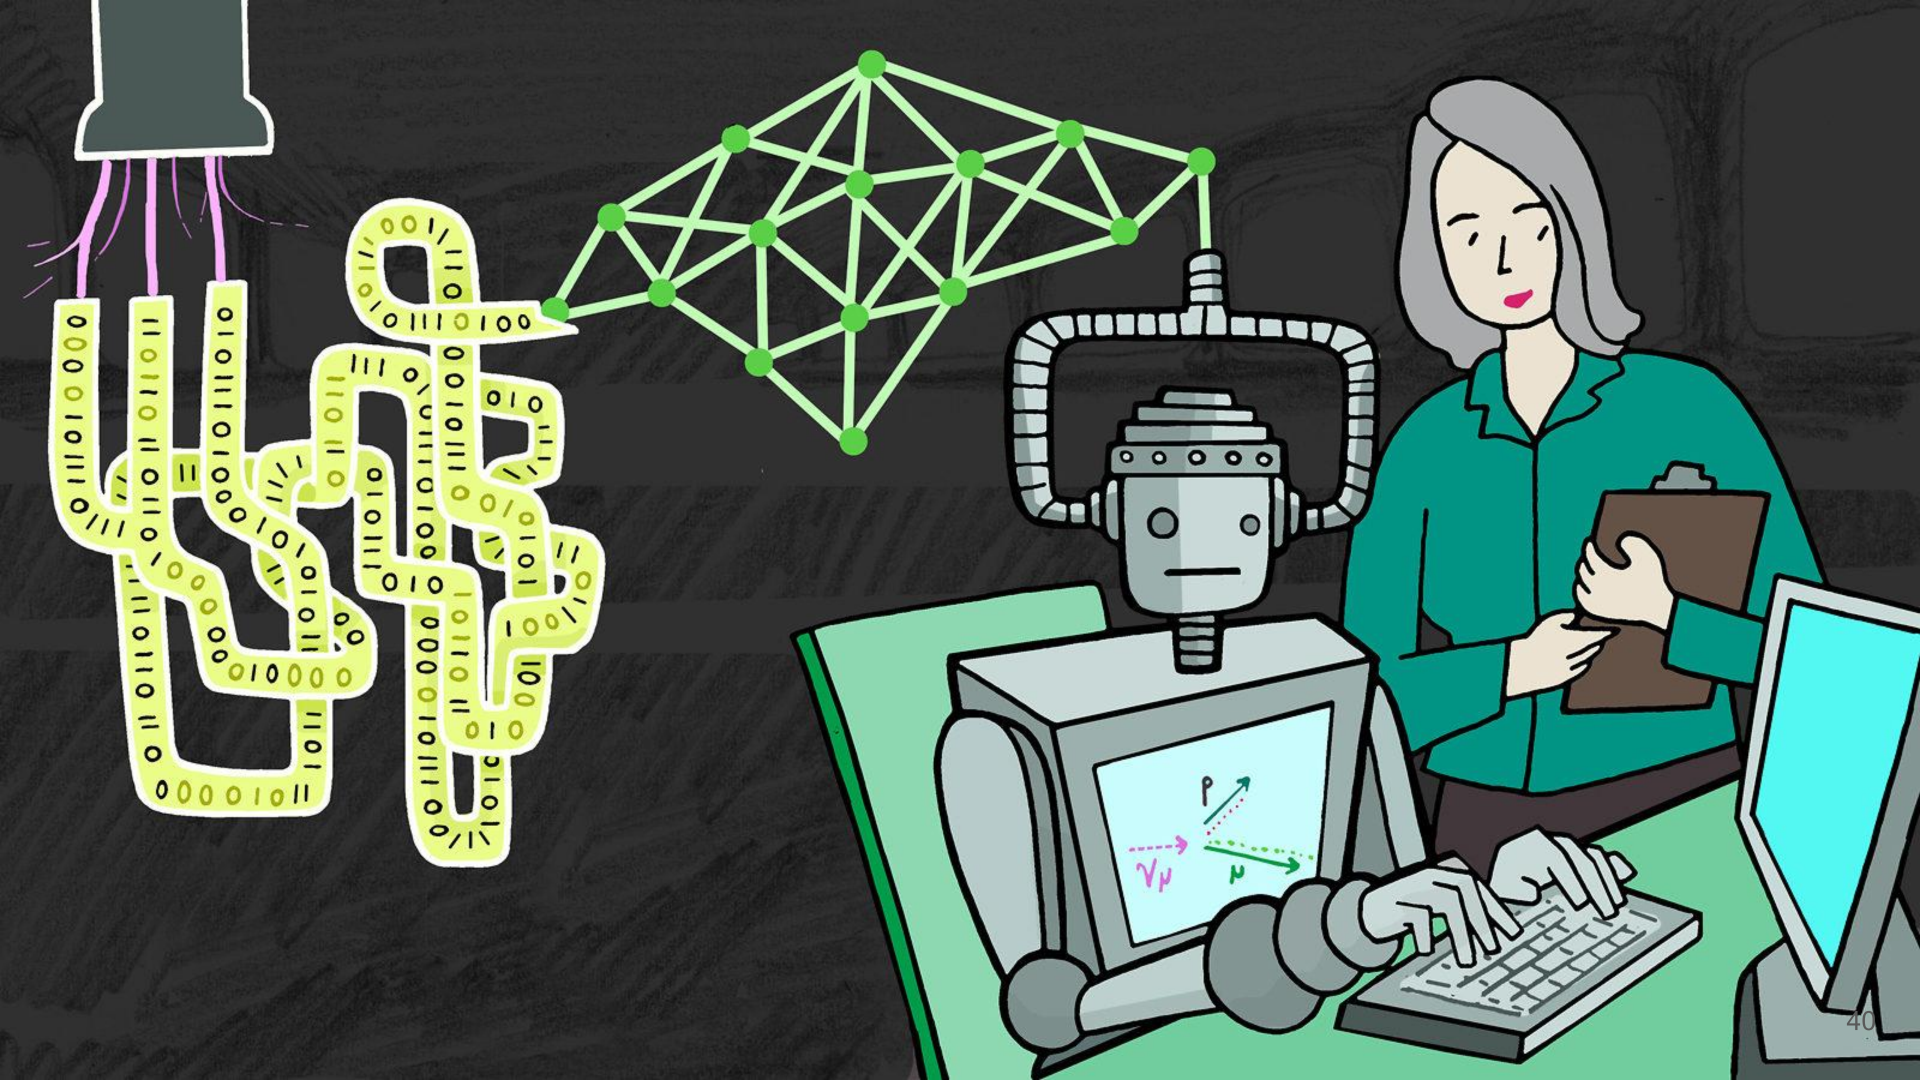
p
ν_μ
μ

SPINE US “Network”

SLAC NATIONAL ACCELERATOR LABORATORY

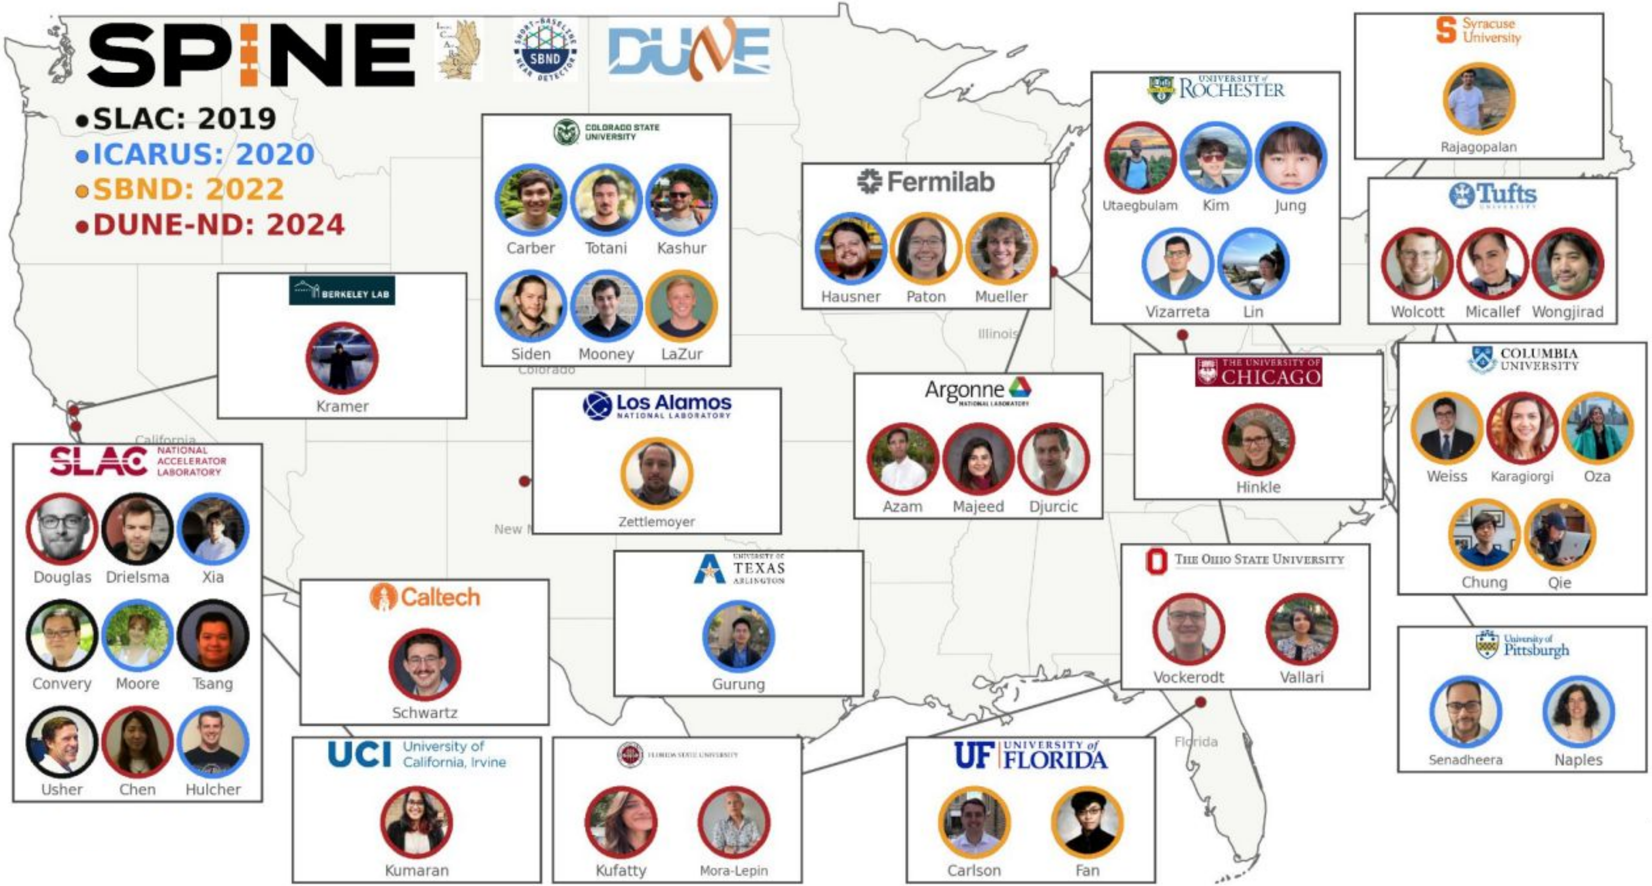

Bonus: Fast surrogate for intra-nuclear modeling

GiBUU final-state interaction

Patrick Tsang
NPML 2025

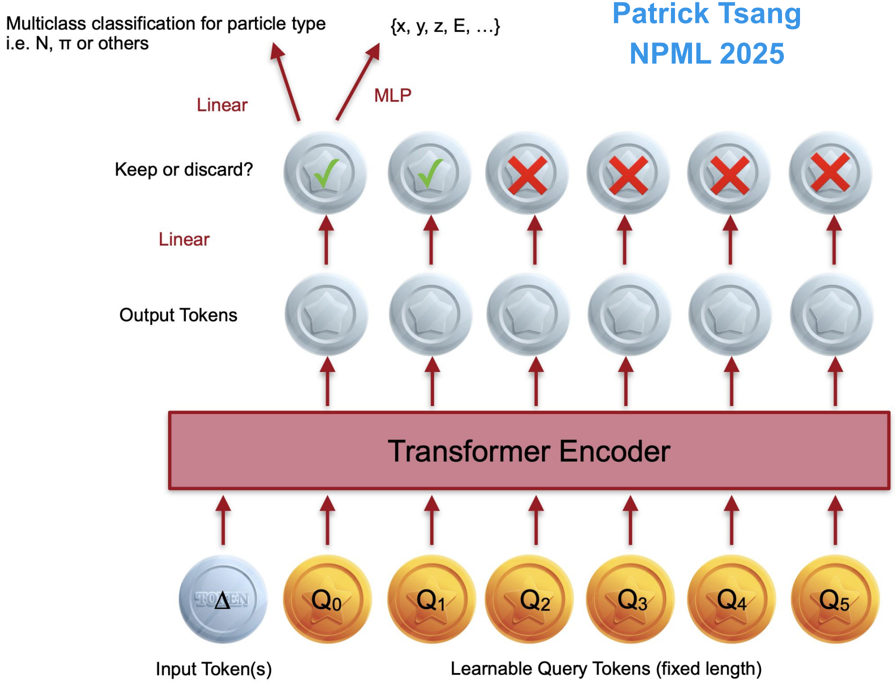

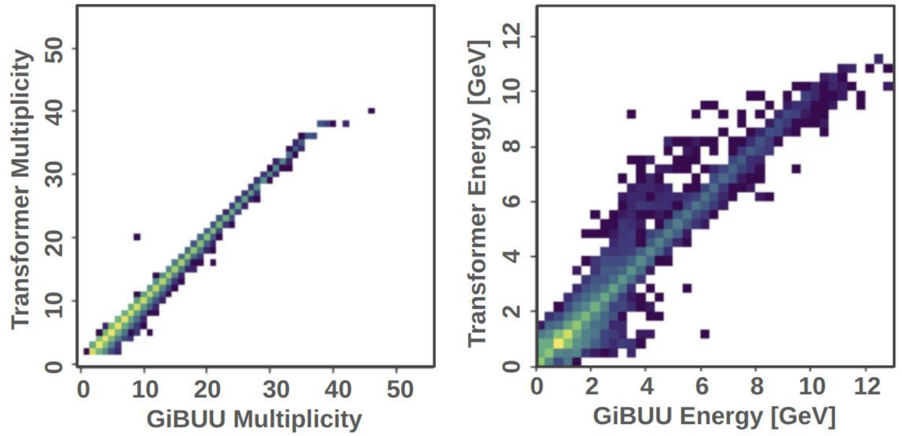

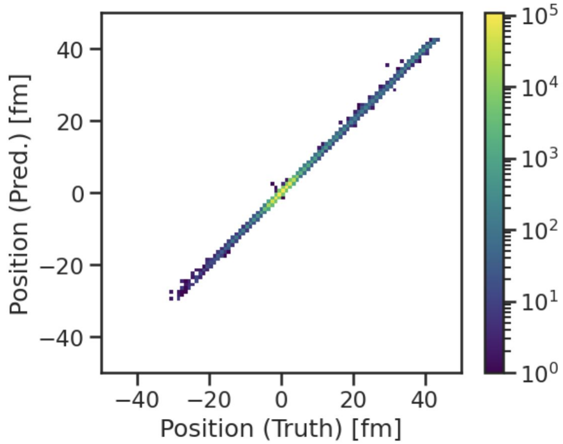

Bonus: Fast surrogate for ν/e-nucleus interactions

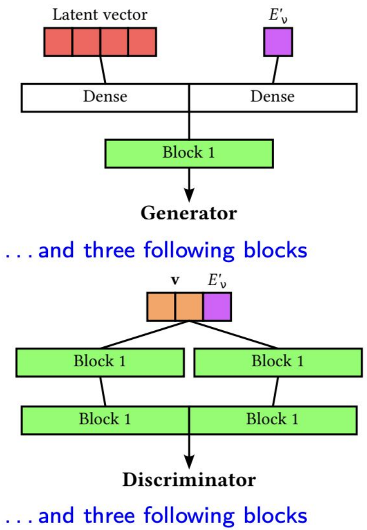

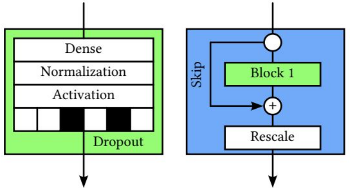

- For each block in the inclusive model, attach the neutrino energy as input.

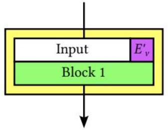

- Add Gaussian noise layers to mitigate mode collapse.

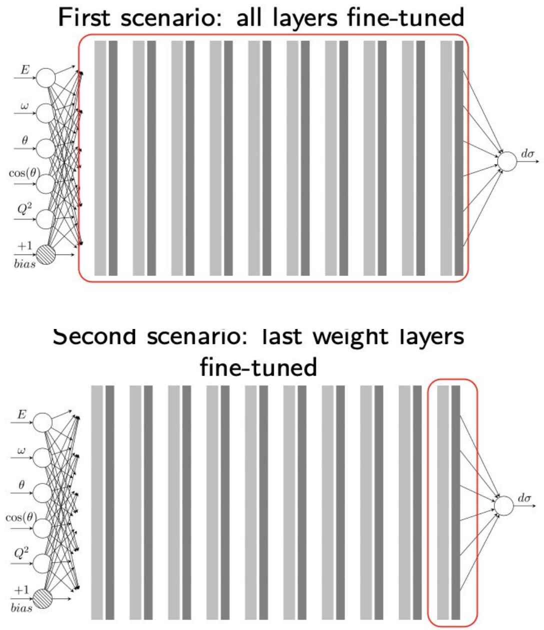

Generating 1 million events on a single-core CPU takes

- approximately 12 minutes with NuWro (full simulation),
- and around 28 seconds with GAN.
- When GAN predictions are made on a GPU (such as RTX 4080), the time is reduced to approximately 3 seconds.

PRD 112 (2025) 013007
arXiv:2508.12987
Krzysztof Graczyk
NPML 2025