

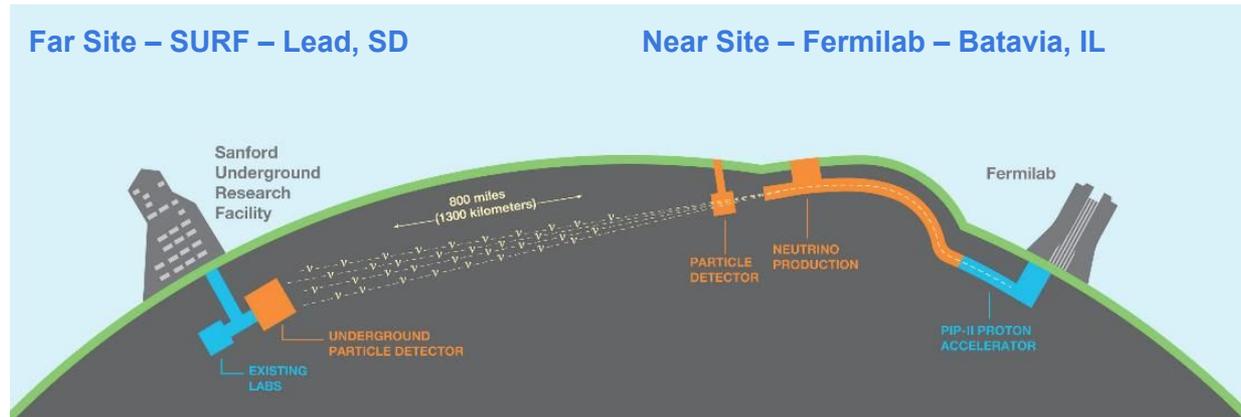
DUNE Collaboration Meeting – Update for NPPS

Feb 25, 2026
Jake Calcutt

Introduction to DUNE

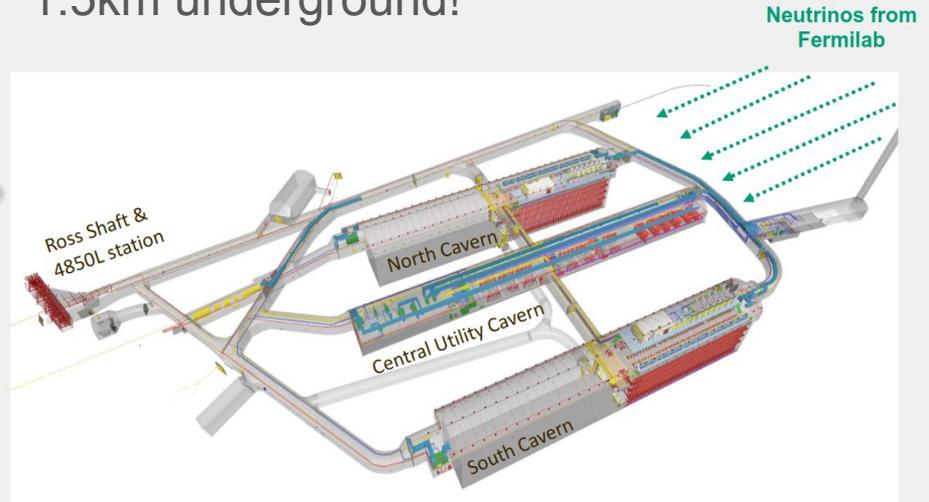
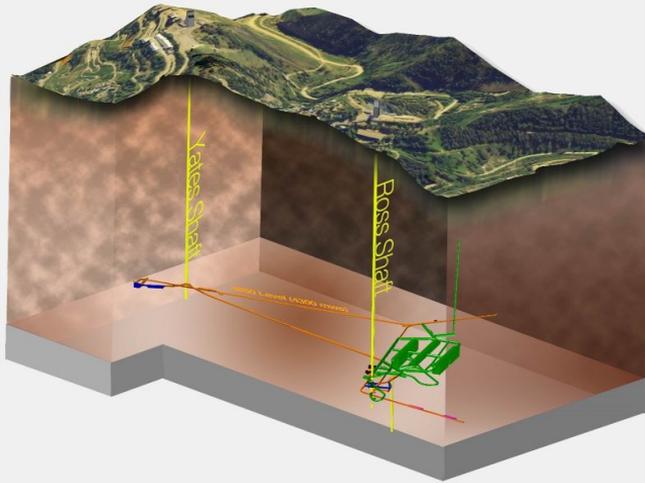
Upcoming next-gen long baseline neutrino oscillation experiment

- Target physics:
 - Oscillations – neutrino mass ordering, CP violation, precision oscillation parameter measurements
 - Supernova detection & Pointing (~5 Degrees)
 - BSM – Nonstandard nu interactions, dark matter, HNL, etc.



LBNF Far Site at SURF (South Dakota)

1.5km underground!



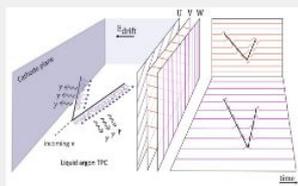
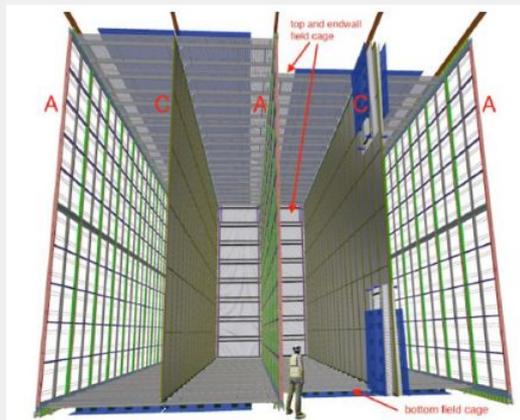
SURF caverns completed in 2024

Construction starting this year: Tentatively Aug 3!



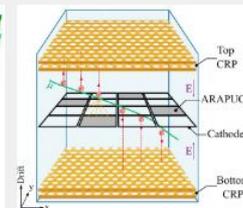
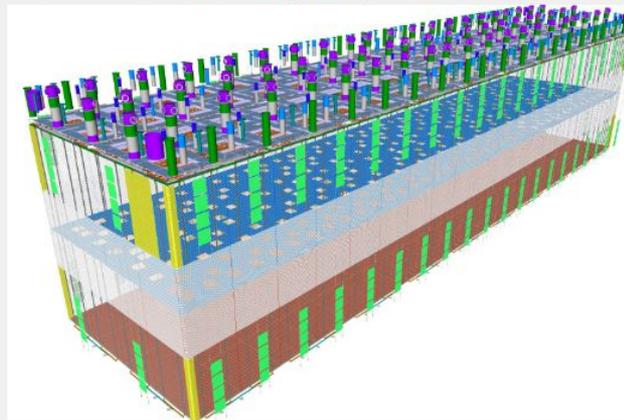
Far Detector Technologies

FD-HD: JINST 15 (2020) 08, T08010



H 12 m x
W 14 m x
L 58 m

FD-VD: JINST 19 (2024) 08, T08004



H 12 m x
W 14 m x
L 60 m

Horizontal drift (FD-HD)

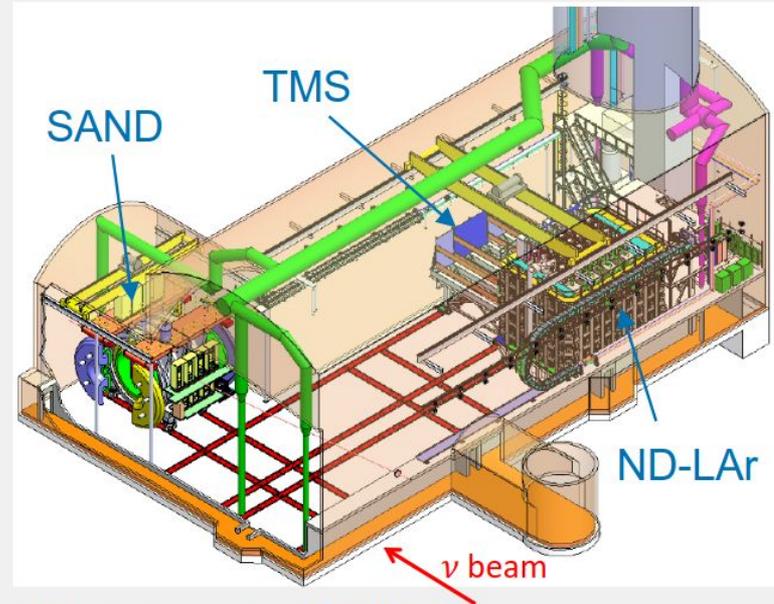
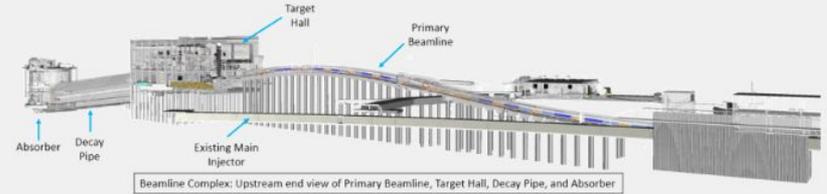
- Four 3.6 m drift regions
- Wire readout plane (APA) technology
- X-ARAPUCA photon detectors integrated in anode planes

Vertical drift (FD-VD)

- Two 6.25 m drift regions and central cathode
- Charge readout plane (CRP) technology
- X-ARAPUCA photon detectors integrated in cathode plane and membrane walls

Near Detector Complex

- Main purpose: enable prediction of Far Detector reconstructed spectra
- Movable detector system (DUNE-PRISM): LArTPC (**ND-LAr**) with muon spectrometer (**TMS**)
 - Off-axis data in different neutrino fluxes constrains energy dependence of neutrino cross sections
 - Same target, same technology → inform predictions of reconstructed E_ν in Far Detector
- On-axis magnetized detector (**SAND**) for beam monitoring and neutrino measurements



CDR: Instruments 5 (2021) 4, 31



CERN Neutrino Platform



ProtoDUNE-VD
(770 ton LAr)

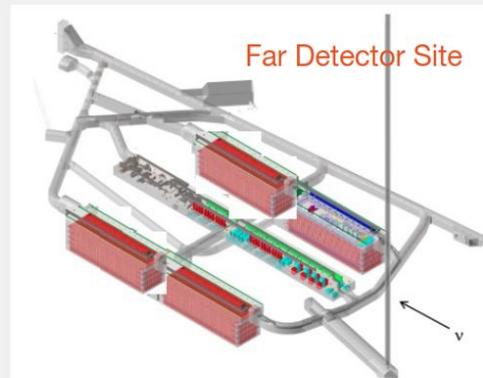
ProtoDUNE-HD
(770 ton LAr)

ProtoDUNE/DUNE ~1/20
Full scale DUNE FD components

Unique cryogenic infrastructure for R&D
on LArTPCs at large-scale

DUNE Phases

- **DUNE Phase I** (2026 start detector installation; early 2030 physics; 2032 beam + ND)
 - Full near + far site facility and infrastructure
 - Two 17 kt LArTPC modules
 - Upgradeable 1.2 MW neutrino beamline
 - Movable LArTPC near detector with muon catcher
 - On-axis near detector
- **DUNE Phase II:**
 - Two additional FD modules (≥ 40 kt fiducial in total)
 - Beamline upgrade to >2 MW (ACE-MIRT)
 - More capable Near Detector (including ND-GAr)



Physics Updates – Submissions/Publications

Publications: submitted since last collaboration meeting

- 2601.05697, Reconstruction of atmospheric neutrinos in DUNE's horizontal-drift far-detector module
- 2511.13462, Measurement of Exclusive π^+ -argon Interactions Using ProtoDUNE-SP
- 2511.11925, First Measurement of π^+ -Ar and p-Ar Total Inelastic Cross Sections in the Sub-GeV Energy Regime with ProtoDUNE-SP Data
- 2510.08380, Identification of low-energy kaons in the ProtoDUNE-SP detector
- 2509.07664, Towards mono-energetic virtual ν beam cross-section measurements: A feasibility study of ν -Ar interaction analysis with DUNE-PRISM
- 2509.07012, Operation of a Modular 3D-Pixelated Liquid Argon Time-Projection Chamber in a Neutrino Beam

Publications: published in 2025

- JINST (2025) 20 P09008, Spatial and Temporal Evaluations of the Liquid Argon Purity in ProtoDUNE-SP
- EPJC 85, 697 (2025), Neutrino interaction vertex reconstruction in DUNE with Pandora deep learning
- Phys. Rev. D 111, 092006 (2025), Supernova Pointing Capabilities of DUNE

My work!

Computing-focused presentations

Framework Development

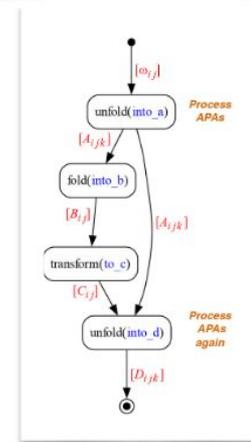
- Undertook a **External Review of Conceptual Design Report**

- 5 HEP framework developers reviewed
- 1.5 day review in Oct 2025 - <https://indico.fnal.gov/event/71180/>
- review was extremely positive -
- final report has been posted to [docdb](#)
- “Its design confers **clear advantages** over existing reconstruction frameworks for neutrino experiments such as DUNE.”
- “The review team has been **very impressed** with the overall design of the new Phlex framework, and commends the Phlex developers for going through a design process that includes gathering requirements and user stories. “
- “The review committee concluded that there are technical risks at the conceptual level. However, **we believe these risks are acceptable, and that they can be mitigated with appropriate considerations.**”



(Parallel, Hierarchical, and layered execution of data)

- mechanisms for keeping memory use bounded
- methods of handling time dependent conditions data
- navigation between child/parent data products when earlier products have been evicted from memory
- High Order Functions with very large numbers of inputs



NPPS Input could be good for this sort of task!

Test using LLMs for adapting algorithms to new framework

Chatbot Summary on Bruce Baller's Lifetime Module

Bot Name	Could get code from GitHub URL	Skeleton or attempt at implementation?	Completed request	Made pointers or copied data	Followed associations	Comments
ChatGPT 5o	yes	skeleton	yes	copied data	yes	Gave me an alternative implementation (some other code)
Copilot m365	yes	skeleton	yes	pointers	no	Too little code to answer some questions
Claude Sonnet 4.0	yes	full	no (and had to be nudged)	copied data	yes	stopped outputting mid-statement. Could not make more even when asked
Claude Sonnet 4.5	yes	full	yes*	pointers, but had to be prodded on clusters (hits ok)	yes	*one of three files was missing after prodding it to use pointers – asking it for the missing file produced it.
Google Gemini	no	full	no*	pointers	yes	* left some code as an “exercise to the reader”

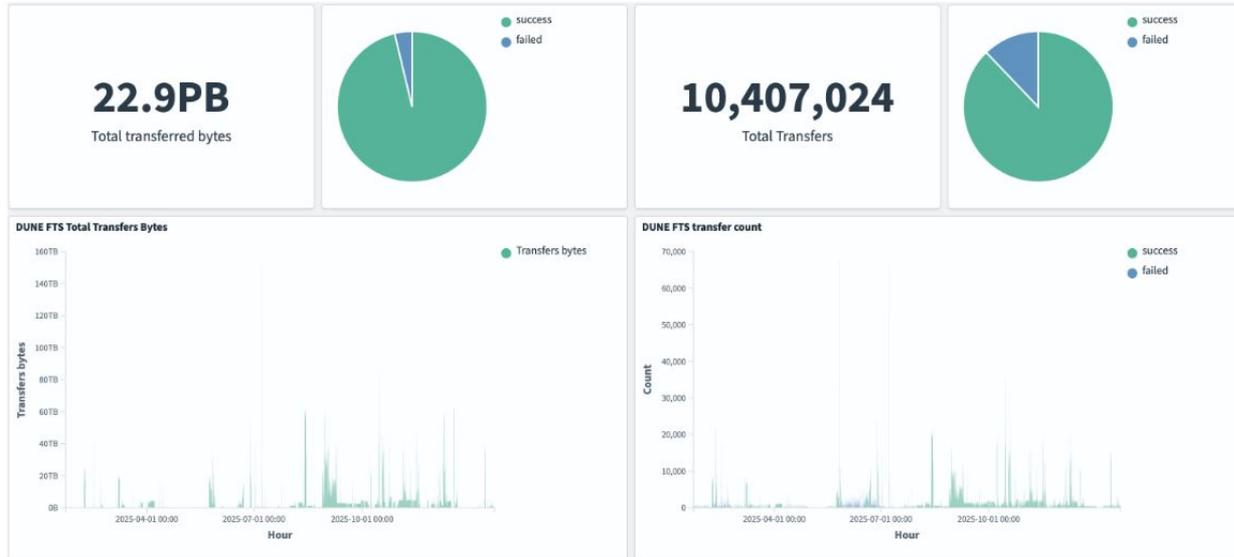
Chatbot Summary on Bruce Baller's Lifetime Module

Bot Name	Could get code from GitHub URL	Skeleton or attempt at implementation?	Completed request	Made pointers or copied data	Followed associations	Comments
openwebui.fnl.gov qwen3-coder:latest (30.5 B)	no (uploaded)	full	yes	pointers	no	Failed to output histograms
openwebui.fnl.gov mixtral:latest (46.7 B)	yes	skeleton	yes	not enough code	no	invented dependence on opDetWaveforms
Perplexity	yes	skeleton	yes	not enough code	no	I asked again for an implementation and got a “you’ll need to fill in” a lot
openwebui.fnl.gov llama3.2:latest (3.2 B)	yes(?)	not enough	yes?	not enough code	no	It thought we were using a timer. Model is way too small.
openwebui:llama3.1:8B	no (uploaded)	skeleton	yes	not enough code	no	got stray things about cluster efficiency and SNR

Plans for 2026

- Data Management
 - Requirements/design/implementation for bulk upload of output files (joint with Workflow)
 - Plans for supernova-format data challenge
 - Data Placement Agent - this will allow us to respond to data set needs more rapidly and in a more automated way
 - clean up CERN EOS and evaluate our policy of raw data access (how does this affect the DAQ if at all?)
- FORM development - continues and it is an open question what library will contain the raw data file interface into Phlex - but continued coordination/discussion needed
- Updates to the Conditions DB for PDVD - improvements/expansion on the schema and logging of additional information

Data transfers in the past year



We managed about 22.1 PB transfers using Rucio and Data Pipeline via FTS in the past year.



Note: source of Rucio scalability crisis is DUNE's current paradigm of individual uploads of data + logs from individual jobs (as opposed to batch uploads from a dropbox)

Scalability studies

- Over the course of 2025, we observed an increasing number of Rucio service slowdowns.
- During these incidents, Rucio server pods experienced:
 - high CPU load
 - Numerous low-level unhandled exceptions
 - Rapid growth in memory usage
- As memory usage continued to rise, affected pods were eventually restarted by Kubernetes health checks.
- Further investigation showed that these events were associated with large numbers of near-simultaneous Rucio uploads coming from jobs which had finished at approximately the same time
 - 500–1000 jobs finishing within a 2–3 minutes window were sufficient to trigger this failure



Scalability studies

Yuyi Guo & Dennis Lee

- We implemented a series of optimizations on the Rucio server side:
 - Increased pod memory: 3GB -> 4GB
 - Increased listen-backlog(wsgi): 100 -> 200
 - Increased httpd timeout: 60s -> 300s (5m)
 - Increased liveness probe time
- On January 15, 2026, we migrated from an IP-based load balancer to an Ingress-based load balancer, resulting in significantly improved load balancing and overall scalability.
- Now an extreme load case of 1600 simultaneous Rucio uploads was nevertheless able to cause some timeouts, although not pod restarts.
- We conclude that changes to the Rucio upload based workflow are still warranted.
- More details and discussions will be in the afternoon session

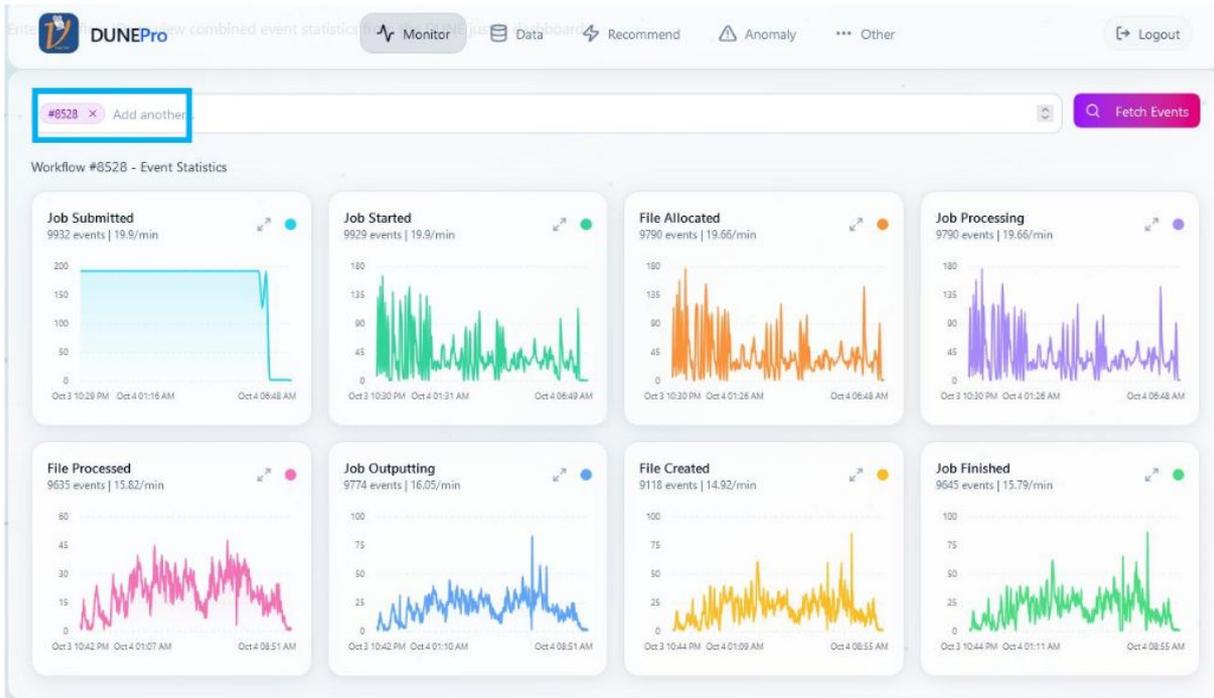
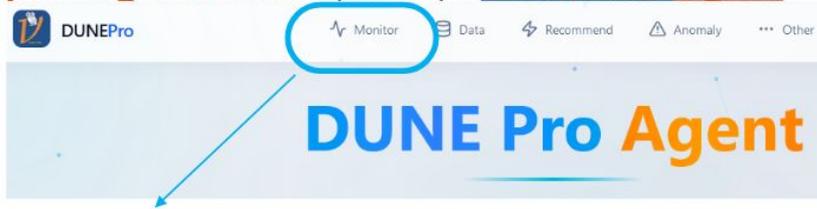


Data Placement Agent

Note: JustIN is our main distributed workflow management system

- There has always been a box in the system diagram most frequently referred to as the “Data Placement Agent” and most commonly understood to be a Data Management responsibility.
- JustIN can send workflows to sites where the data is closest but it doesn’t have capability to identify when data could be replicated to improve the efficiency of a workflow
- Hit this over the summer—all the “reco1” of the far detector LE Monte Carlo was sitting at Fermilab but the reco2 was running slowly because there were few slots available in the US.
- Goal is to design an agent which may or may not be AI-based to identify opportunities where copying a dataset would speed up a workflow
- Some interest from US and UK on this. No final personnel assignment made as yet.

▪ new tool for reporting available (RICE) : **DUNE Pro Agent**



Monitor workflow execution

- *Detect issues*
- *Prepare summary plots (CPU, RAM,)*

justIN events to OpenSearch

- justIN's central services log events when "something happens"
 - Listed here https://dunejustin.fnal.gov/docs/event_types.md
 - They are put in the MySQL/mariadb database
- Chris Brew and now Jyoti Biswal at RAL have produced a containerised agent to send these events to the DUNE OpenSearch managed by Wenlong at Edinburgh
- This is working at the prototype level
- It produces **extremely** useful, long lived monitoring of workflows, jobs, site and storages successes and failures
- Workflow needs to fill out some unfilled fields in some of the events
 - Already have some feature requests for metrics (CPU time, file size, ...) in addition to the current discrete fields (Workflow ID, Site Name, ...)
 - justIN 2.0 will create events for failed upload attempts too

Explore MCP in DUNE Rucio

- From the Rucio workshop, Belle II implemented a PoC of a Rucio MCP server that can be used by shifters or end-users to query for Rucio (check rules, list files, etc) [link](#)
- We'd like to try a similar Proof of Concept, to explore the possibility of using MCP in DUNE Rucio, opensearch monitoring, JustIN...
- Locally, Edinburgh University provides an OpenAI API key we can use as a LLM host for trail use
- Then I built a MCP server/client that can talk to DUNE Rucio, opensearch monitoring (coding by ChatGPT...)
- The following examples give us potential use cases



Explore MCP in DUNE Rucio

```
You: count total number of rules created by dunepr0
Assistant: The total number of rules created by the account "dunepr0" is 116,754.
You: how many files in hd-protodune:hd-protodune_trigprim_26300
Assistant: The dataset "hd-protodune_trigprim_26300" in the scope "hd-protodune" contains a total of 93 files.
You: which rules related to hd-protodune:hd-protodune_trigprim_26300
Assistant: There is one rule related to the dataset "hd-protodune_trigprim_26300" under the scope "hd-protodune":

- **Rule ID:** 22e01b9dd30d4cdfac2ed44b4890f0eb
- **State:** OK
- **RSE Expression:** RAL-PP
- **Account:** dunepr0
- **Locks OK Count:** 93
- **Grouping:** DATASET
- **Priority:** 3
- **Activity:** User Subscriptions
- **Created At:** 2025-12-17T20:23:13
- **Updated At:** 2025-12-19T21:18:37

This rule is associated with the RAL-PP RSE and has a total of 93 locks that are in an OK state.
```



Explore MCP in DUNE Rucio

```
You: list rses in a table, grouped by country, where "Availability Read" is True
Assistant: Here's a table of storage elements (RSEs) with "Availability Read" set to True, grouped by country:
```

Country	RSE Name
CERN	CERN_PDUNE_CASTOR CERN_PDUNE_EOS DUNE_CERN_EOS
CZ	PRAGUE
ES	DUNE_ES_PIC
FR	DUNE_FR_CCIN2P3_DISK DUNE_FR_CCIN2P3_TAPE
IN	DUNE_IN_TIER
IT	DUNE_IT_INFN_CNAP
NL	NIKHEF SURFSARA
UK	DUNE_UK_LANCASTER_CEPH QMUL RAL_ANTARES RAL_ECHO RAL_PP
US	DUNE_US_BNL_SOCC DUNE_US_FNAL_DISK_STAGE FNAL_DCACHE FNAL_DCACHE_PERSISTENT FNAL_DCACHE_STAGING FNAL_MISC_TAPE T3_US_NERSC



Explore MCP in DUNE Rucio

```
You: which dataset hd-protodune:np04hd_tp_run026300_0088_tpwriter_tpswriter_20240520T105244.hdf5 belongs to?
Assistant: The dataset 'hd-protodune:np04hd_tp_run026300_0088_tpwriter_tpswriter_20240520T105244.hdf5' belongs to the following datasets:

1. 'hd-protodune_26300'
2. 'hd-protodune_trigprim_26300'
You: where is hd-protodune:np04hd_tp_run026300_0088_tpwriter_tpswriter_20240520T105244.hdf5, give me the url with root://
Assistant: The file 'hd-protodune:np04hd_tp_run026300_0088_tpwriter_tpswriter_20240520T105244.hdf5' is available at the following locations with the 'root://' protocol:

1. **CERN_PDUNE_CASTOR**
- 'root://eosctapublic.cern.ch:1094//eos/ctapublic/archive/neutplatform/protodune/rawdata/np04//hd-protodune/trigprim/2024/detector/cosmics/None/00/02/63/00/np04hd_tp_run026300_0088_tpwriter_tpswriter_20240520T105244.hdf5'

2. **FNAL_DCACHE**
- 'root://fncdcaoor.fnal.gov:1094/pnfs/fnal.gov/usr/dune/tape_backed/dunepr0//hd-protodune/trigprim/2024/detector/cosmics/None/00/02/63/00/np04hd_tp_run026300_0088_tpwriter_tpswriter_20240520T105244.hdf5'

3. **RAL_PP**
- 'root://mover.pp.rl.ac.uk:1094/pnfs/pp.rl.ac.uk/data/dune/hd-protodune/92/51/np04hd_tp_run026300_0088_tpwriter_tpswriter_20240520T105244.hdf5'
```

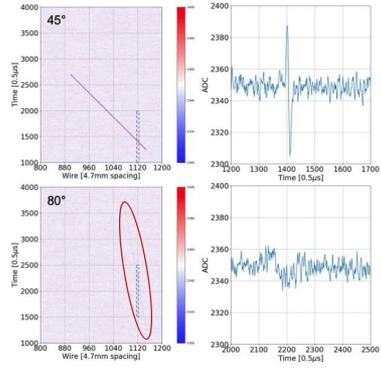


Introduction

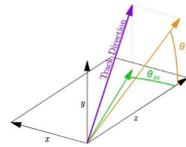
- Fermilab has been migrating to the Spack package management system as part of the transition to Alma Linux 9 (AL9).
- Although Scientific Linux 7 (SL7) has passed its end-of-life, collaborators still primarily utilize the UPS package management system inside SL7 Apptainer containers.
- Since the release of Spack 1.0 in July, DUNE and the FNAL spack migration team have been working towards full adoption of Spack for publishing our software ecosystem.

Physics/Reco x ML

DNN ROI motivation

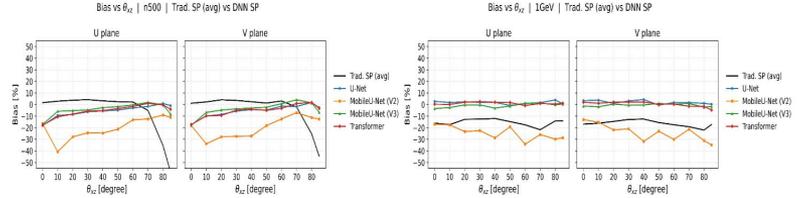


- Object with large θ_{xz} : Prolonged topology
 - Track traveling normal to the wire planes
 - Signals will have low amplitude
- Prolonged topologies lead to poor S/N ratio in induction wire planes
- Object with small θ_{xz} : Isochronous topology
 - Track traveling parallel to the wire planes
 - Signal will occur in multiple channels at the same time



X axis - electric field direction
Y axis - vertical upward
Z axis - beam direction
arXiv: [2007.06722](https://arxiv.org/abs/2007.06722)

Recap - Performance evaluation with different DNN models



Single track

Single shower

For single track:

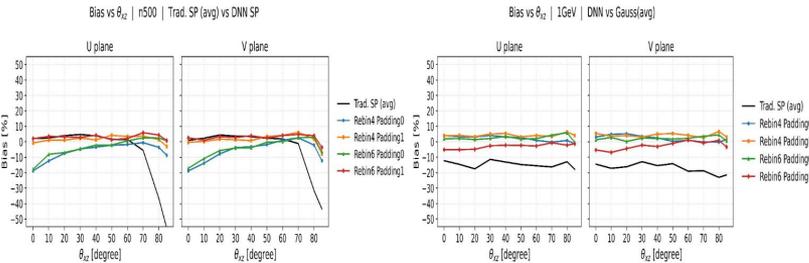
- DNN ROI demonstrated better charge reconstruction performance at high angles
- But at low angles, DNN ROI performs worse than the Trad. ROI

For single shower:

- Three DNN models showed similar performance across the θ_{xz}

[Previous talk](#)

Performance evaluation with Rebin/Padding

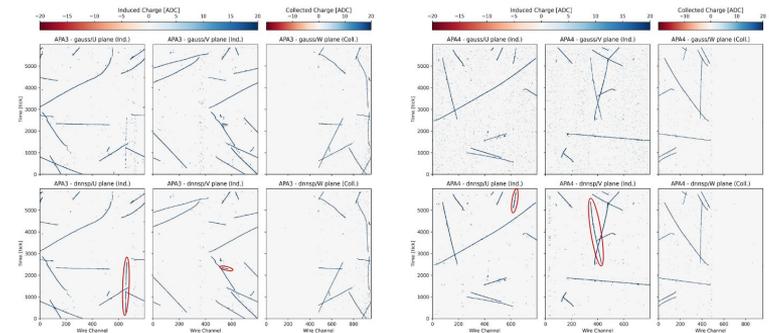


Single track evaluation

Single shower evaluation

- Rebin factor: the image resolution along the time axis
- Implemented padding successfully reduced bias of isochronous track cases
- The best performance is from the model trained with rebin factor = 4, with padding

DNN ROI with real PD-HD data

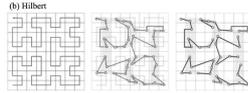
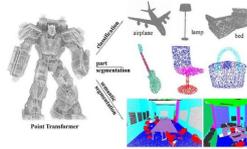


PD-HD data (028588-4562)

PD-HD data (027673-1)

Point Set Transformer

- Directly take point clouds as input, naturally suited for sparse LArTPC data.
- Applies permutation-invariant self-attention between points in point cloud, appropriate for unordered point sets.
- Complexity grows as N^2 , so attention operation usually limited to kNN (or point serialization method in Point Transformer V3).



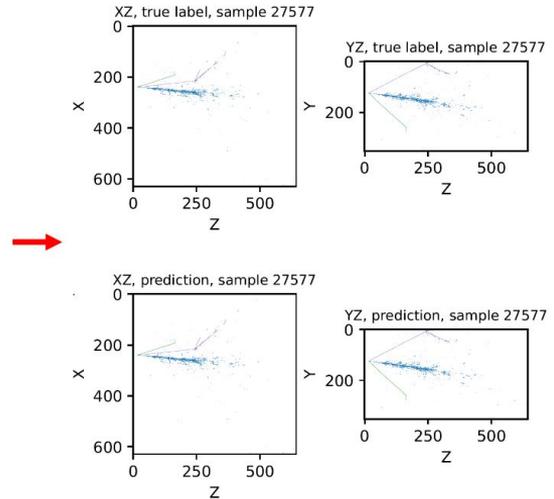
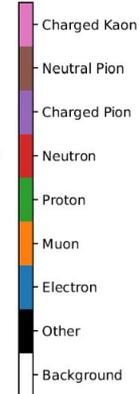
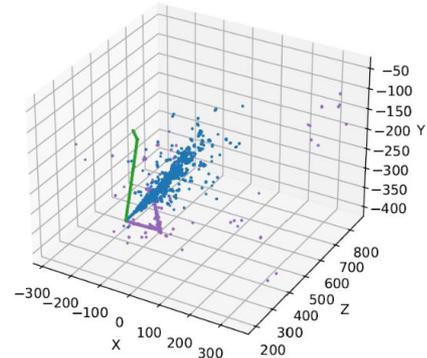
"Point Transformer" IEEE / CVF ICCV 2021

"Point Transformer V3: Simpler, Faster, Stronger" IEEE / CVF CVPR 2024

Point Set Transformer for Semantic Segmentation

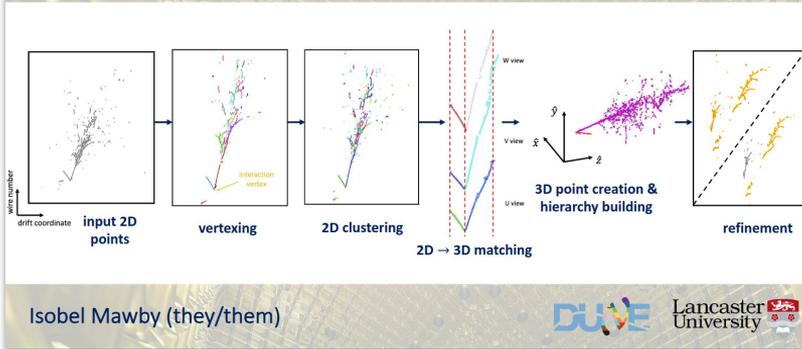
Event Display Examples

Sample 321, 3D view

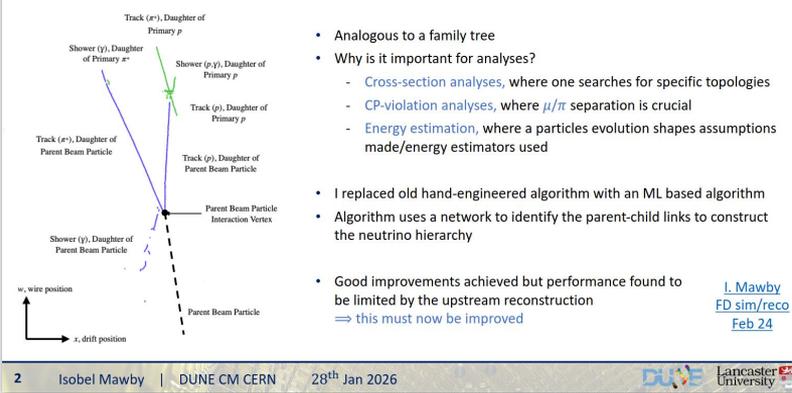


Pandora: 'mainline' reco framework w/ mix of trad. and ML routines

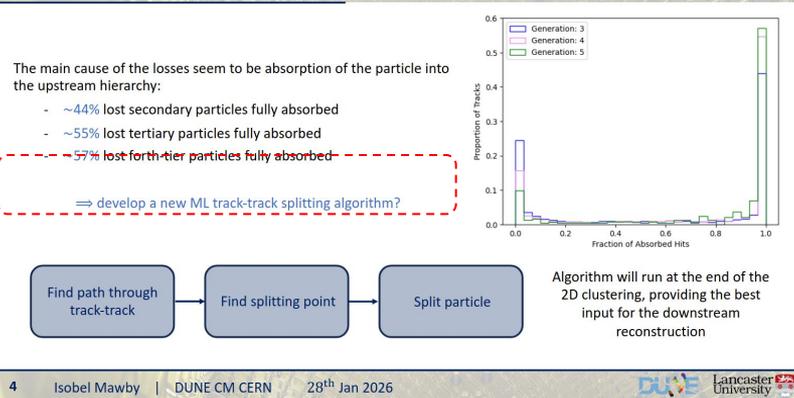
Hierarchy Building in Pandora for DUNE: Cluster Splitting



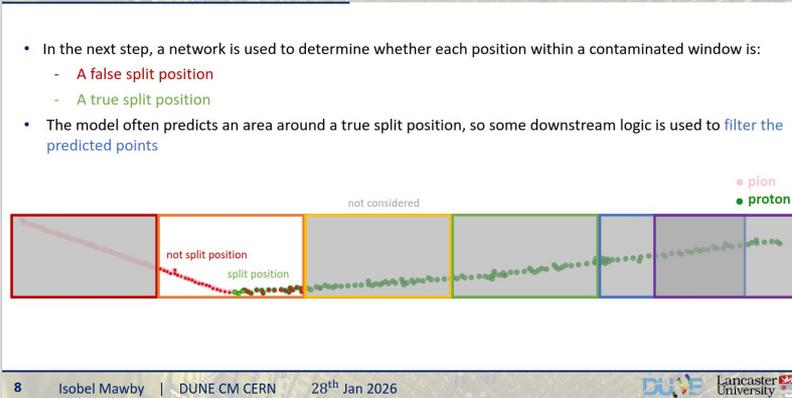
Neutrino Hierarchy Construction



Why are our particles not being reconstructed?



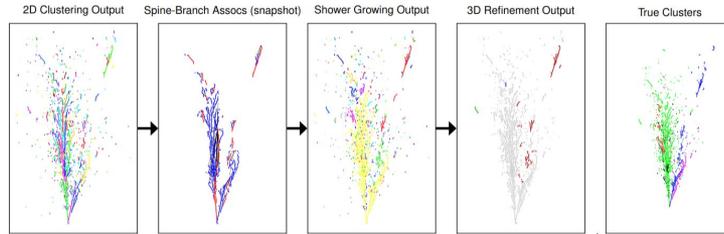
Cluster Splitting Algorithm: 3. Where is the Split?



Shower clustering w/ Transformer model

Shower Growing

- Shower growing task: merge clusters to form complete showers from the output of 2D clustering
- Current hand-engineered algorithms fail frequently for showers with angular contamination from other particles



Shower Growing Task

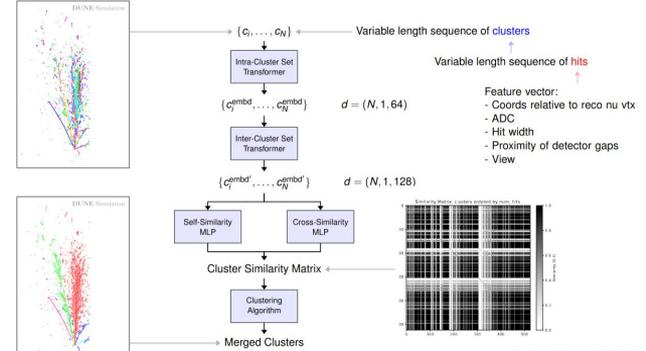
Alex Wilkinson

Transformer-Driven Shower Clustering in Pandora

2 / 14

The Model

Pandora 2D clustering output



Alex Wilkinson

Transformer-Driven Shower Clustering in Pandora

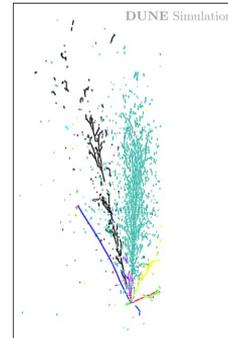
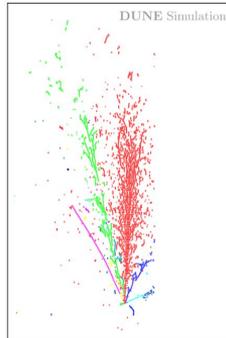
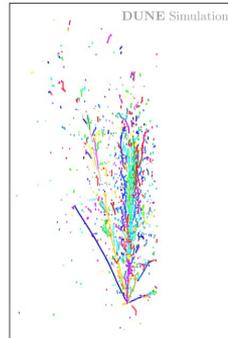
4 / 14

Examples

Input

DL Clustering

True Clustering



Alex Wilkinson

Transformer-Driven Shower Clustering in Pandora

8 / 14

Motivation & Goals

Anomaly?

Beyond Standard Model (BSM) Physics

Signatures from exotic particles, including boosted dark matter, millicharged particles, magnetic monopoles, etc.

Supernova Burst Neutrinos

Rare but known physics events requiring prompt identification for multi-messenger astronomy

Detector Malfunctions

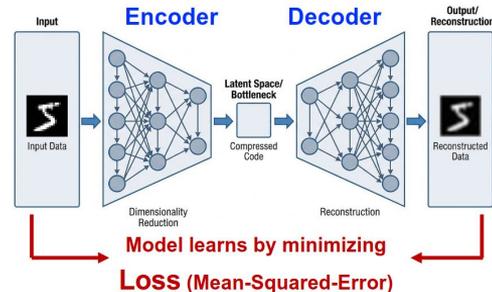
Unexpected behavior in detector components requiring intervention

Why Real-time on Raw Data?

- ❖ Detect rare physics events before standard reconstruction
- ❖ Reduce data volume by pre-filtering at front end
→ Accelerate discovery
- ❖ Enable prompt SNB triggers for multi-messenger astronomy
- ❖ Monitor detector in real-time for quality assurance

Autoencoder Model Illustration

- **Encoder**: learn to compress input data into compact latent representation
- **Decoder**: reconstruct original input from latent space



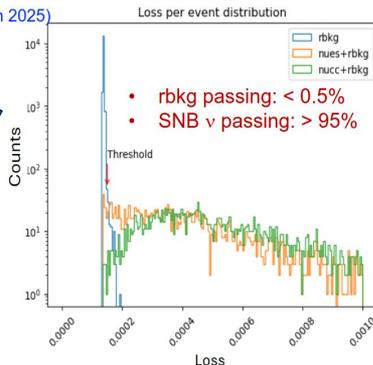
Proof of Concept Results (Summer 2025)

Eric Novello (Fermilab SULI intern 2025)

- The final model was evaluated using evaluation data.

Evaluation Data

- 16k rbkg events (held-out, not used for training)
- 1,920 SNB ν CC events + rbkg
- 1,920 SNB ν ES events + rbkg



Next Steps & Implementation Path

1. Enhanced Training Data

- Include **solar ν** as nominal bkg
- Add **cosmic μ** events to training dataset
- Develop detector malfunction simulation scenarios

2. BSM Physics Evaluation

- Test on **Boosted DM** signatures
- Evaluate millicharged particle sensitivity
- Assess dark-sector tridents detection
- Test magnetic monopoles and atm. LLPs

3. FPGA Deployment

- Optimize model (pruning, quantization)
- Benchmark latency, power consumption, efficiency
- Implement using HLS4ML or CGRA4ML

4. Operational Integration

- Retrain on real data from pDUNE/ND2 \times 2
- Adapt to continuous data stream handling
- Extend to U/V wire planes if needed
- Develop agentic AI shift/monitoring system

Wire-Cell (WC) Event Reconstruction & Porting Progress

Part 1: Sim-NF-SigProc → Part 2: Img-Clus-QLMatching-nu-select → Part 3: 3D PatRec

<p>TPC simulation</p> <p>noise filtering</p> <p>signal processing</p>	<p>3D imaging</p> <p>clustering</p> <p>charge-light matching</p>	<p>3D trajectory & dQ/dx fitting</p> <p>cosmic muon taggers</p>	<p>multi-track fitting</p> <p>3D PatRec</p> <p>neutrino taggers</p> <p>DL-3D vertexing</p> <p>RNN EE</p>
---	--	---	--

[JINST 12 P08003 \(2017\)](#)
[JINST 13 P07006 \(2018\)](#)
[JINST 13 P07007 \(2018\)](#)
[JINST 16 P01036 \(2020\)](#)

[JINST 13 P05032 \(2018\)](#)
[JINST 16 P06043 \(2021\)](#)

[Phys. Rev. Applied 5, 064071 \(2021\)](#)

[JINST 17 P01037 \(2022\)](#)

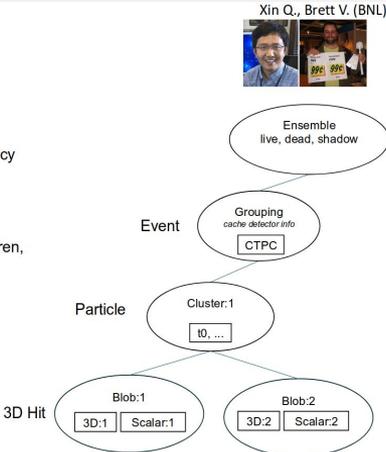
Brookhaven National Laboratory

Porting: Data Structure

- To use uboone clustering at SBND and DUNE:
- Expand to **multi-APA** and **wrapped wire** detectors
 - All blobs now carry the APA/face information
 - All algorithms rewrite accordingly
 - Major refactor of code
 - Adapt to wire-cell-toolkit framework for higher efficiency
 - e.g., memory efficiency

New data structure **PointCloudTree** to reduce memory usage

- Cluster::3D() ~ a view from 3D PointCloud in all children, no copy needed
- Use **Facade** to realize traditional operations
- Utilities to convert PointCloudTree to *ITensorSet*
 - A general interface
- This structure is capable of incorporating downstream PatRec algorithms

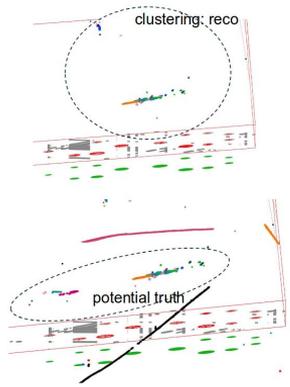


Main Wire-Cell developments last year

- WC Part 1 Maintenance and Development** – Wenqiang G. (formerly at BNL), Brett V. (BNL), Joseph Z. (FNAL), Gray P. (FNAL), Xuyang N. (BNL), Jay J. (BNL), Haiwang Y. (BNL)
 - Bending cathode; YZ-patched sim; resampling; paddings in SigProc; [NF/SigProc for PD](#); hydra for FD; etc
 - [wire-cell-toolkit](#), [larwirecell](#), [dunereco/DUNEWireCell](#)
- WC Part 2 & 3 Development and Validation** – Xin Q. (BNL), Brett V., Lynn T. (UChicago), Ewerton (formerly at LSU), Prabhjot S. (LSU), Hanyu W. (LSU), Haiwang Y.
 - Multi-apa generalization
 - Porting the uboone alg. to the general purposed Wire-Cell Toolkit
- AI/ML developments**
 - DNN-ROI
 - Deployment at SBN - Moon J. (UChicago), Gray P. (FNAL), Avinay B. (UChicago), Haiwang Y.
 - Computing resource reduction - Hokyong N. (Chung-Ang University), Dikshant S. (UCI), Edgar E. (formerly at UCI), Yiwen X. (UCI), Haiwang Y., etc.
 - AI infrastructure for Wire-Cell
 - Triton for Wire-Cell - Inference as a Service – Mike W. (FNAL), Meghna B. (FNAL)
 - Truth labelling, File format - Xuyang N., Jay J., Haiwang Y.
 - DL-clustering – Avinay B., Lynn T., Haiwang Y.

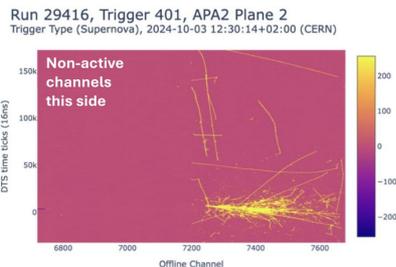
Enhancing the neutrino clustering using deep learning

- One problem identified in uboone
 - Clustering of NCpi0 parts ...
 - Multiple activities, not be able to clustered together
 - Clustering problem (before Q-L matching)
- <https://www.phy.bnl.gov/twister/bee/set/d620d2d4-0608-47b9-a5aa-169b9d743c4d/event/291/>
- Separate gammas are not clustered together for NCpi0 → important for photon related analysis



Triggering for Neutrinos/BSM at ProtoDUNE

- Major challenges for the trigger:
 - Surface detectors – **huge cosmic background** (DAQ can write **max 5/10 events per second**)
 - SPS has a long **4.8s beam spill** and **15s cycle time**
- Need **online self-trigger** algorithm
- Used existing trigger algorithm that selects **high energy events parallel to the cathode**
- Developed ML-based online algorithm

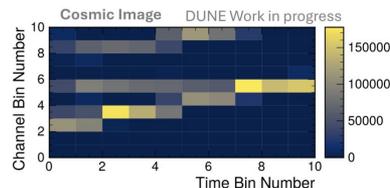
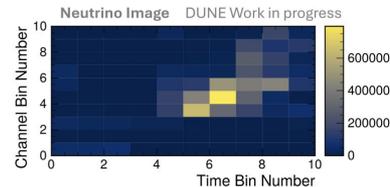


10. 27/01/2026

Ciaran Hasnip | DUNE DRA

Choice of model: XGBoost

- XGBoost**: framework for building and training **gradient boosted decision tree models (GBDT)**
- Important consideration - **speed!**
- ProtoDUNE DAQ has strict latency requirements – **NNs slow** on CPUs
- Trained XGBoost model is **exported** (via [Treetite](#) → [TL2cgen](#)) to **compatible C++**
- XGBoost model compiled as an object in the DUNE DAQ software for fast **sub-microsecond inference**

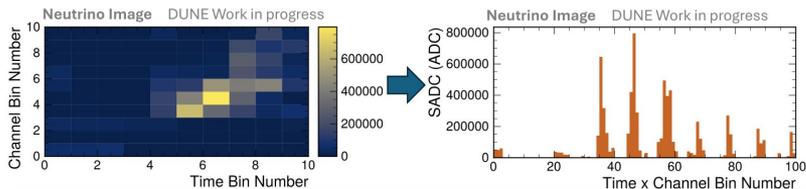


16. 27/01/2026

Ciaran Hasnip | DUNE DRA

Training XGBoost Model

- XGBoost **classifies** between low-res images that are **neutrino or cosmic-like**
- Model sees is an array of numbers – **1D projection** of the 2D image
- Generate 0000s of simulated cosmic and neutrino samples at ProtoDUNE-HD for training
- Simple XGBoost configuration: 50 tree, max depth = 4, learning rate = 0.5

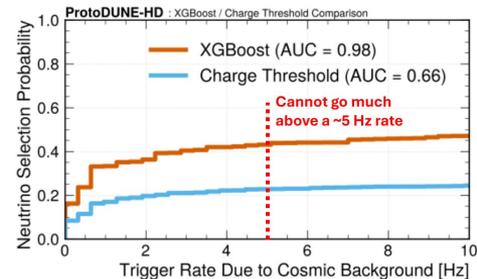


17. 27/01/2026

Ciaran Hasnip | DUNE DRA

XGBoost Model Performance

- XGBoost model **better at separating signal from background** than original charge-threshold TA algorithm
- Performance **gap grows as more cosmic background events are allowed through trigger** (i.e. higher trigger rate)



19. 27/01/2026

Ciaran Hasnip | DUNE DRA

SPINE: End-to-end ML Reconstruction

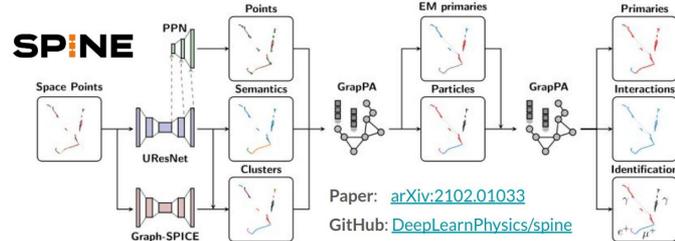
- Developed first at DUNE Near Detector(s)
 - Particularly busy events (see below) – upstream interactions with cavern walls
- Adoption at Far Detector / ProtoDUNEs ongoing

Scalable Particle Imaging with Neural Embeddings



SPINE is an end-to-end optimizable neural-network cascade:

- Voxel semantic classification, point identification (Sparse-CNN: [UResNet+PPN](#), L. Dominé)
- Dense clustering (Smart DBSCAN, Sparse-CNN: [Graph-SPICE](#), D.H. Koh)
- Particle aggregation, shower primary identification (GNN: [GrapPA-Track/Shower](#))
- Interaction aggregation, particle identification, primary identification (GNN: [GrapPA-Interaction](#))



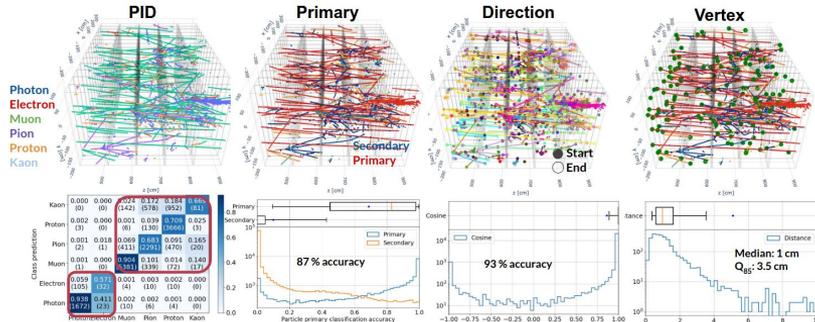
SPINE for DUNE ND-LAr, F. Drielsma

2

μ P4.1 Reconstruction Performance



As presented at the last CM, things look mostly sound at the low level



SPINE for DUNE ND-LAr and its prototypes, F. Drielsma

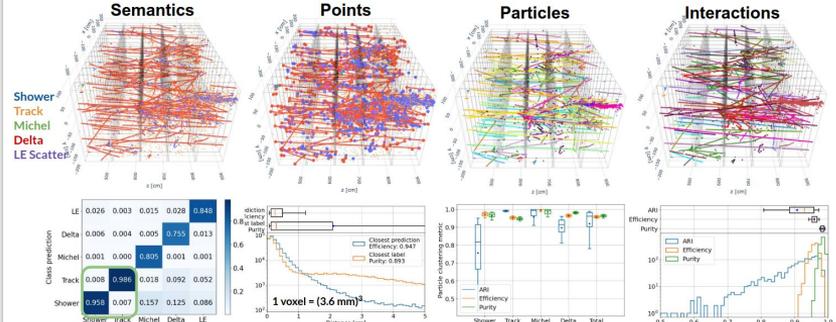
Metrics evaluated for fiducial interactions

44

μ P4.1 Reconstruction Performance



As presented at the last CM, things look mostly sound at the low level



SPINE for DUNE ND-LAr and its prototypes, F. Drielsma

Metrics evaluated for all interactions (with rock)

43

Miscellaneous

Opticks for DUNE Far Det. Simulation

- I remember Gabor was interested in working on something similar for DUNE? What is the status of that?

Particle Simulations

The path of particles through matter is simulated with GEANT

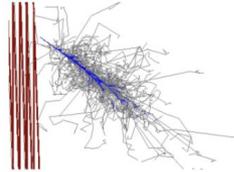
Step-by-step transport: Geant4 simulates particle transport by breaking down the path into small, discrete steps. The path is not a smooth curve but a series of points.

Trajectory: A trajectory is a collection of all the steps and step points for a single particle's path.

The light yield: Scintillation light per MeV deposited is ~25000-50000 optical photons! Tracking that amount of particles is extremely costly (computing)

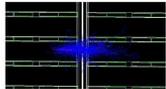
Workaround solutions:

- Look-up tables
- Semi-analytic calculations
- Generative models



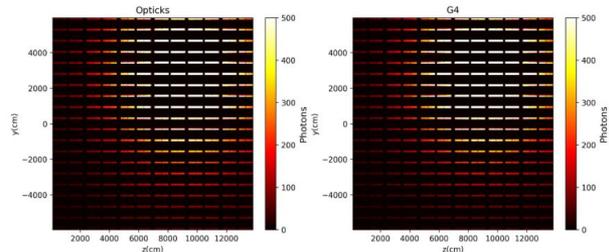
2

Accelerating Optical Simulation in DUNE



dune10kt_v5_refactored_1x2x6_nowires.gdml

3 GeV electron
Produced parallel to z-axis
480 Arapucas



8

Accelerating Optical Simulation

In ray tracing or optical simulation the most important question is, where this ray (photon) "hit" something

Each ray(photon) is independent of each other but can be traced in parallel in a GPU

GEANT4 (CPU)

↓
Generate scintillation photons one-by-one
↓
Trace each sequentially
↓
~10⁴ photons/sec



Opticks (GPU)

↓
Generate scintillation photons in bulk
↓
Trace millions in parallel
↓
~10⁷ photons/sec

6

Software 'framework' for one ND component

- Different ND component technologies led to branching in software/ecosystems
- Has anyone heard of μ fw?

New sw framework: μ fw

- framework with minimal features as a bridge between sandreco and DUNE reconstruction framework
- C++ application engine with plugin algorithms
 - Load only the modules you use
 - Simpler building, dependencies are somewhat isolated
 - .json configuration to decide what to run and how to configure
 - Streamlined ROOT I/O: Automatic dictionaries and trees. Users only see their structs
- sandreco with μ fw is now sandreco-experimental

[Latest presentation](#)

N. Tosi

```
{
  "ufw" : {
    "ufw-loglevel" : "debug",
    "ufw-basepath" : "/usr/local/share/sandreco/data",
    "ufw-ldpath" : ["/usr/local/lib64"]
  },
  "globals" : {
    "sand::geomanager" : { "geometry" : "test/grain.edepsim.root" },
    "sand::grain::geant_run_manager" : {}
  },
  "contexts" : {
    "keys" : 100,
    "locals" : {
      "sand::edep_reader" : { "uri" : "test/grain.edepsim.root" }
    }
  },
  "run" : [
    {
      "sand::grain::optical_simulation" : {
        "geometry" : "geometries/grain/masks-3mm/main.gdml",
        "energy_split_threshold" : 100
      },
      "reqs" : {},
      "prods" : { "hits" : "grain_hits" }
    },
    {
      "sand::grain::detector_response_fast" : {
        "simpl_active_size" : 3.0,
        "simpl_border" : 0.1,
        "pde" : 0.2,
        "reqs" : { "hits" : "grain_hits" },
        "prods" : { "digi" : "grain_digits" }
    },
    {
      "sand::grain::spill_slicer" : {
        "slice_times" : [0.0, 20000.0]
      },
      "reqs" : { "digi" : "grain_digits" },
      "prods" : { "images" : "grain_images" }
    },
    {
      "sand::root::tree_streamer" : {
        "uri" : "test/images.root",
        "tree" : "cameras"
      },
      "write" : { "grain_images" : "sand::grain::images" }
    }
  ]
}
```

The Genesis Mission

- We are generally still waiting for more clarity
- On January 16th DOE announced a Request for Information (RFI)
 - You can find the statement here: <https://www.energy.gov/science/articles/department-energy-seeks-input-advancing-ai-science-and-engineering-workforce>

The Department of Energy (DOE) today announced a Request for Information (RFI) to solicit public and private sector input on strategies for meeting the technical challenges of the [Genesis Mission](#). It also seeks input on developing a skilled American workforce to advance artificial intelligence (AI) in science and engineering.

Deadline: March 8th(!)

The Genesis Mission will mobilize DOE's National Laboratories, industry, and universities to harness the nation's leading capabilities in high performance computing, next-generation quantum computers, and AI to revolutionize science innovation in this nation.

- The DUNE Spokes have asked us to develop a DUNE-specific whitepaper
 - Will use your inputs and could be considered as an RFI