DEEP LEARNING FOR PLENOPTIC NEUTRINO EVENT RECONSTRUCTION

Dr. Saúl Alonso-Monsalve

Neutrino Physics and Machine learning (NPML)

ETH Zürich

16 June 2026

Scintillator-based tracker detectors

T2K SuperFGD - *JINST 13 P02006 (2018)*, *arXiv:2603.14921*

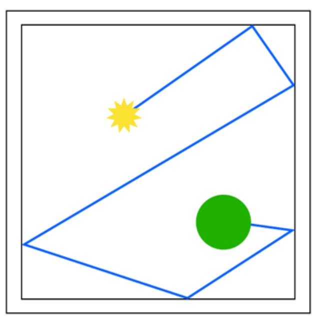

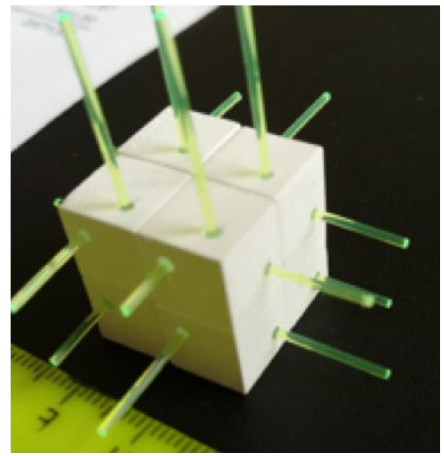

Optically separated segmentation

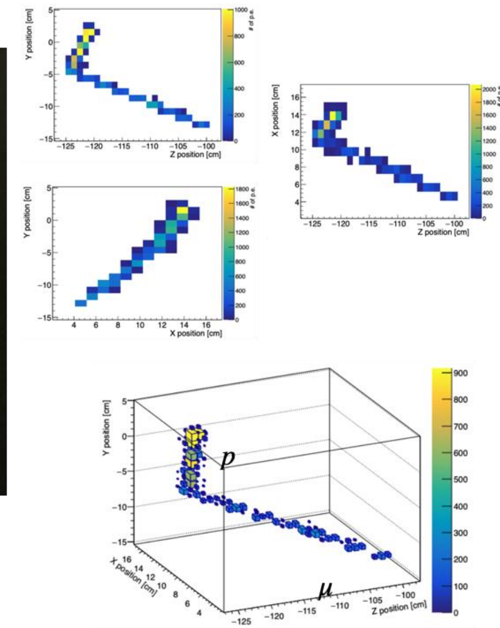

Directional information from individual scintillation photons is not retained.

Spatial information in large transparent scintillators

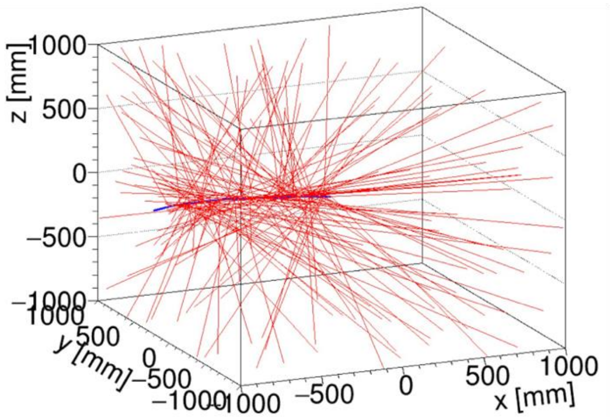

Long attenuation lengths, from metres to tens of metres.

It is commonly assumed that unsegmented, highly transparent scintillator is not ideal for tracking.

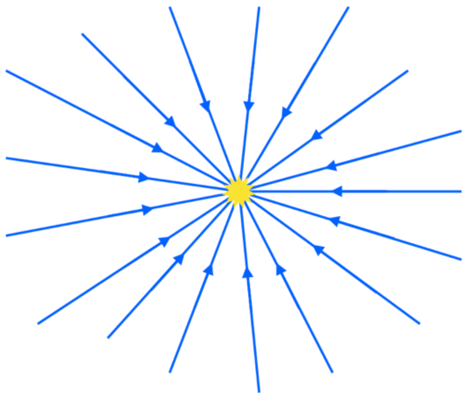

Can we trace individual scintillation photons back to their emission points?

Solution: plenoptic imaging

Fast 3D imaging using a plenoptic camera instrumented with a SPAD array sensor → tracking without segmentation

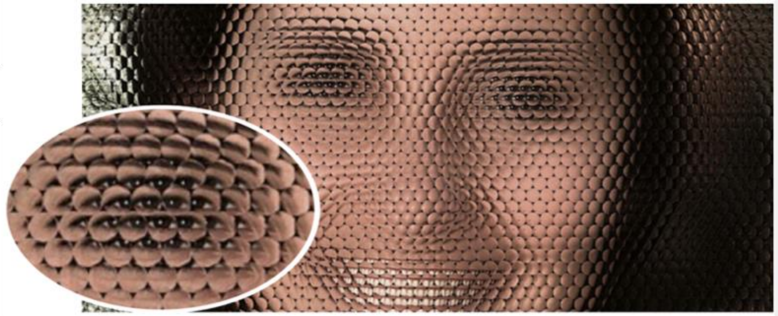

Image from Raytrix Gmbh

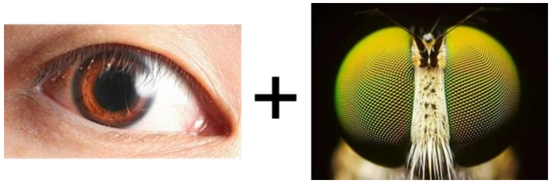

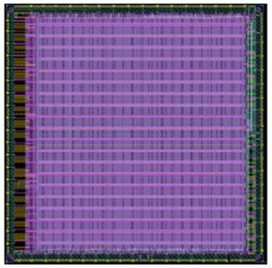

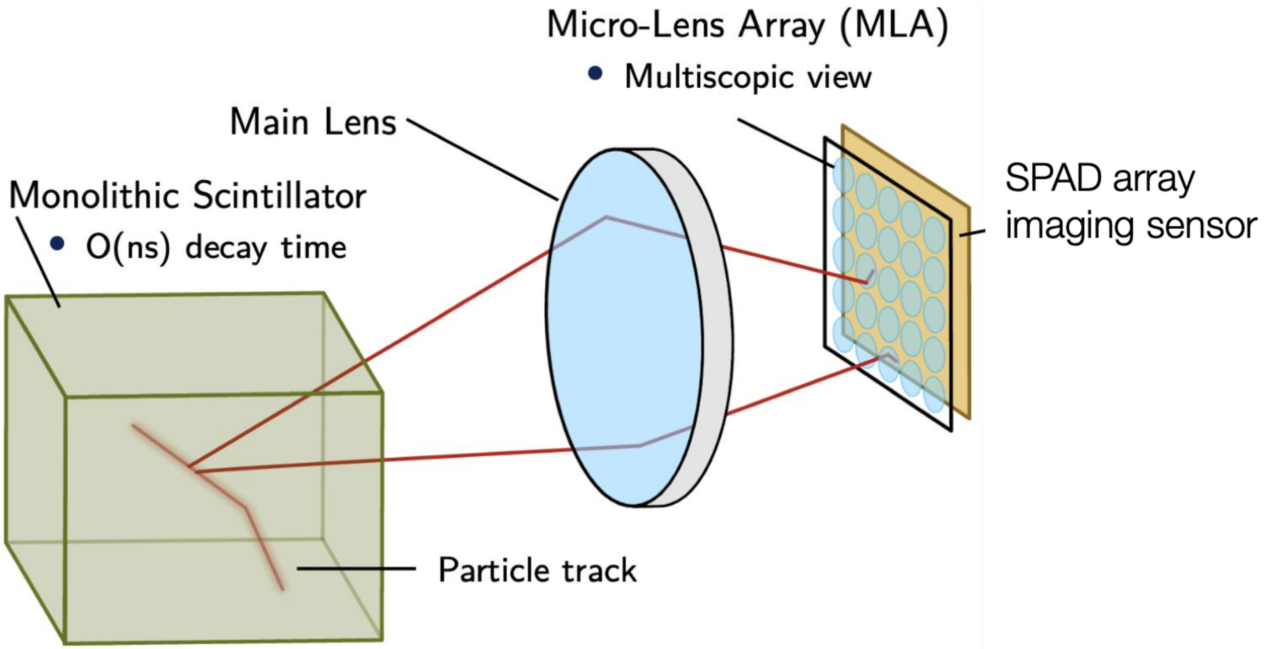

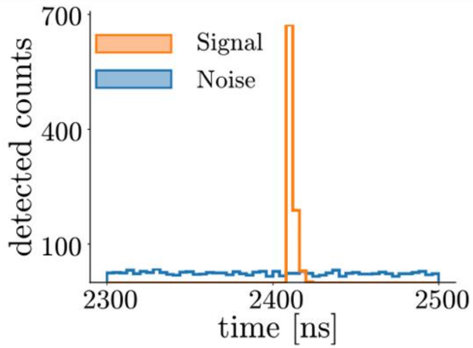

The PLATON detector concept

Nature Communications (2026) - doi.org/10.1038/s41467-026-70918-x.

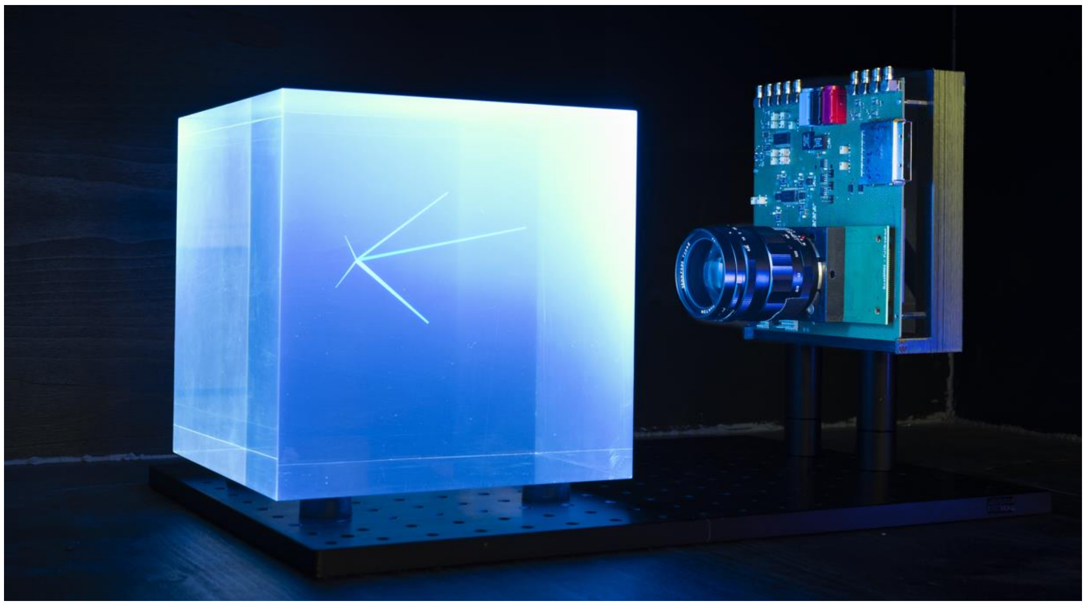

ETH zürich

Swiss National Science Foundation

EPFL

R&D funded by the SNSF (Prof. Sgalaberna, PCEFP2_203261), collaboration with Prof. Charbon's group at EPFL.

Featured by ETH Zürich *Partnerships* and *D-PHYS*.

The PLATON detector prototype

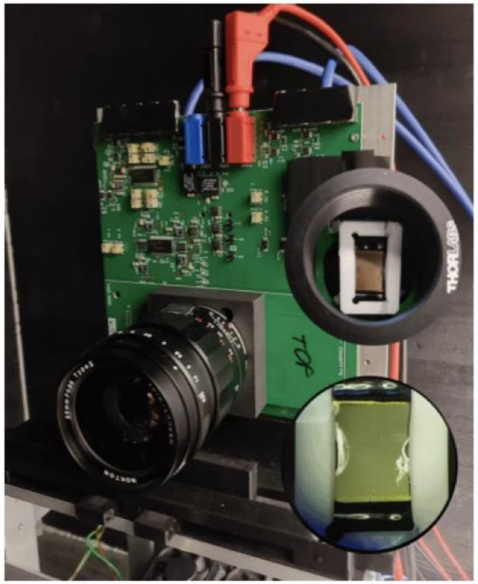

Optics specifications	
Lens	Nokton II 25 mm
f/#	$\sim$2.4
Lateral Resolution	$\sim$1 mm
Depth Resolution	$\sim$5 mm
Depth of Field	$\sim$15 cm

MLA designed and assembled by Raytrix GmbH

SwissSPAD2 Sensor specifications	
Sensor size	512×512 pixel
Pixel pitch	16.38 μm
Photo-detection Efficiency (PDE)	$\sim$5 %
Dark Count Rate (DCR)	0.26 cps/μm^2
Timing	10 ns – 10 μs gated

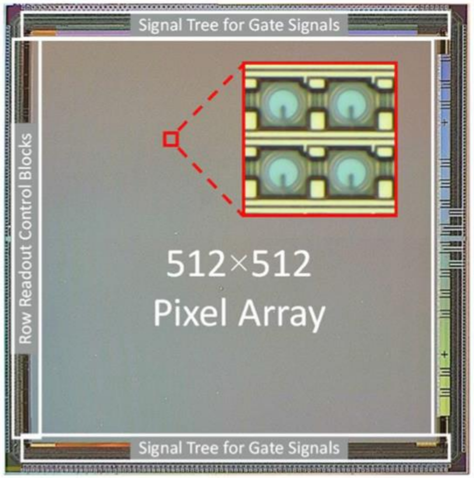

SwissSPAD2 developed by Prof. Charbon's group at EPFL.

[IEEE Journal of Selected Topics in Quantum Electronics 25, 1, 1-12 (2019)]

doi.org/10.1038/s41467-026-70918-x

3D reconstruction in PLATON

- Developed a novel (non-ML) ray-tracing method (Till Dieminger).
  - Tested on ^{90}Sr electrons and “tracks” from two-photon absorption laser.
  - Depth (lateral) resolution: 3 (0.9) mm.
  - Very good correlation between data and Monte Carlo (MC).

- Next challenge: scale the reconstruction to complex neutrino interactions.

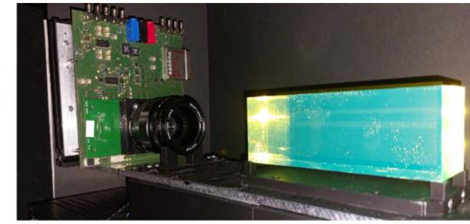

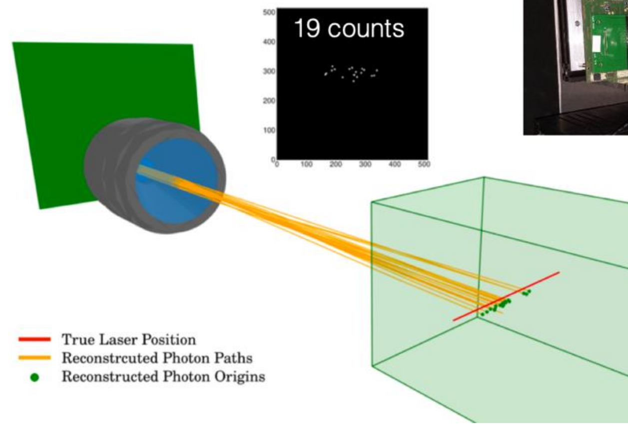

Event

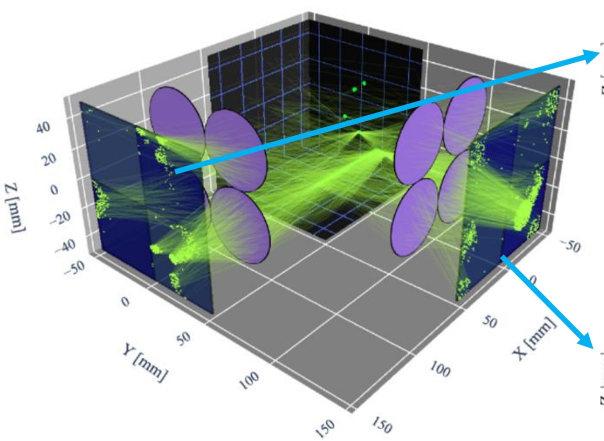

Detected photons

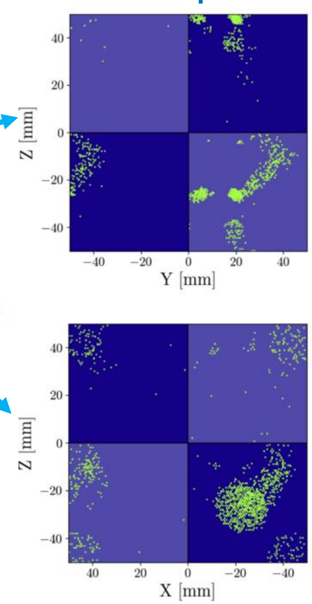

Algorithm

3D reconstruction

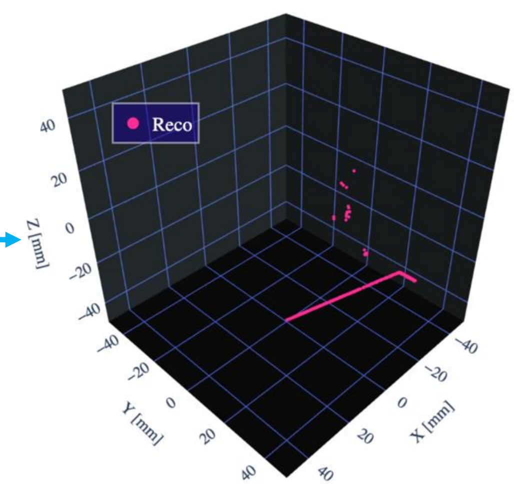

Classical reconstruction

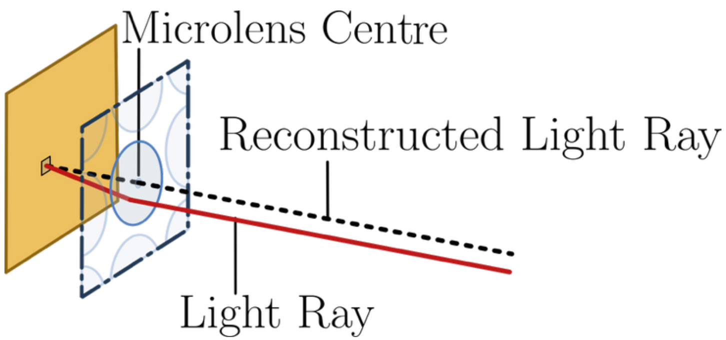

1. Chief-ray method

Trace each pixel through the microlens centre into object space.

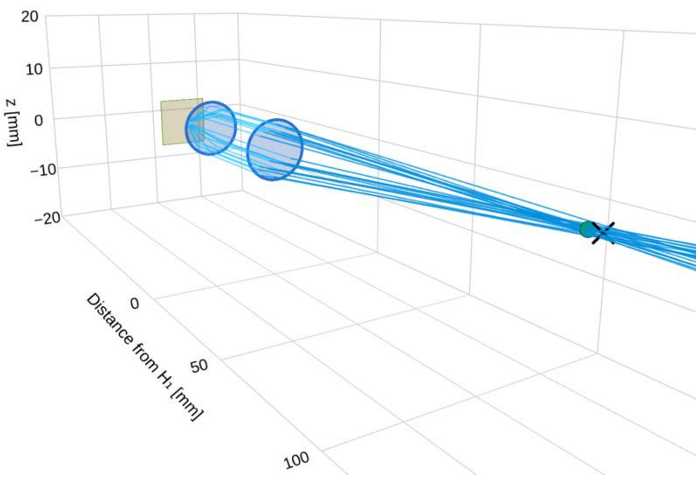

2. Build point cloud

Merge rays from the two orthogonal views.
Find mutual closest-approach pairs.
Midpoint is added to point cloud.

3. RANSAC track fitting

Iteratively draw a line between random points, score by number of close-approaching rays.

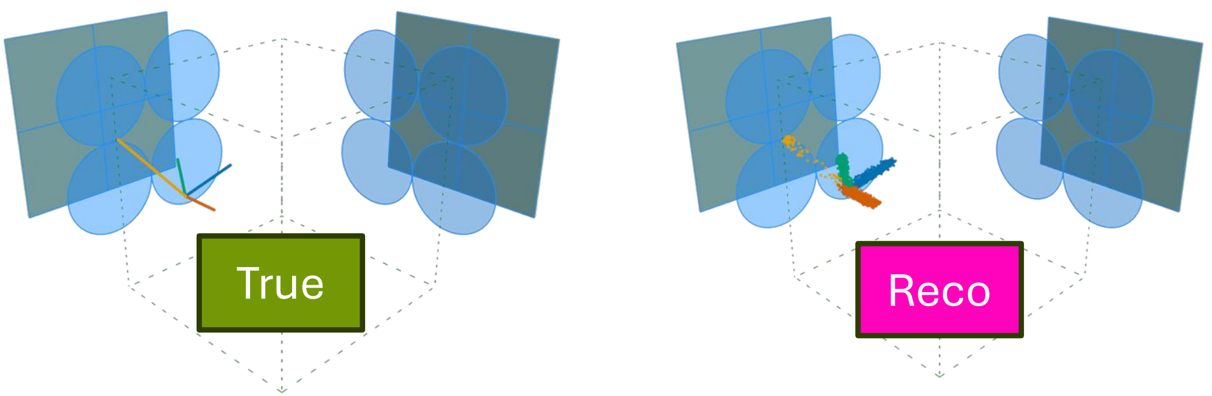

Time consuming!!

PLATON challenge: multiple tracks

- SPAD images: noisy pixels (dark counts) and overlapping photon patterns make reconstruction harder.

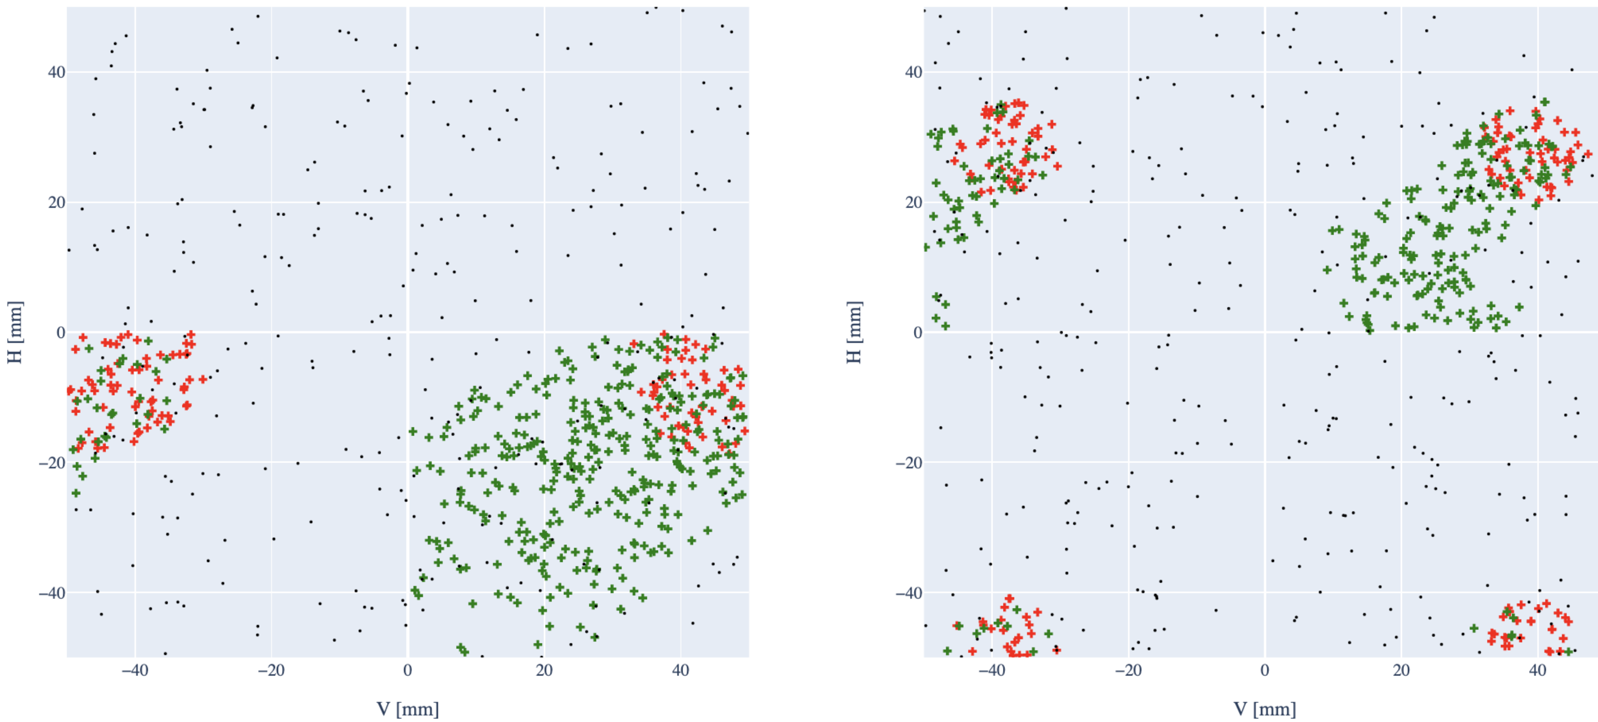

Red = photons from particle A.
Green = photons from particle B.
Black = dark counts.

doi.org/10.1038/s41467-026-70918-x

Deep learning reconstruction in PLATON

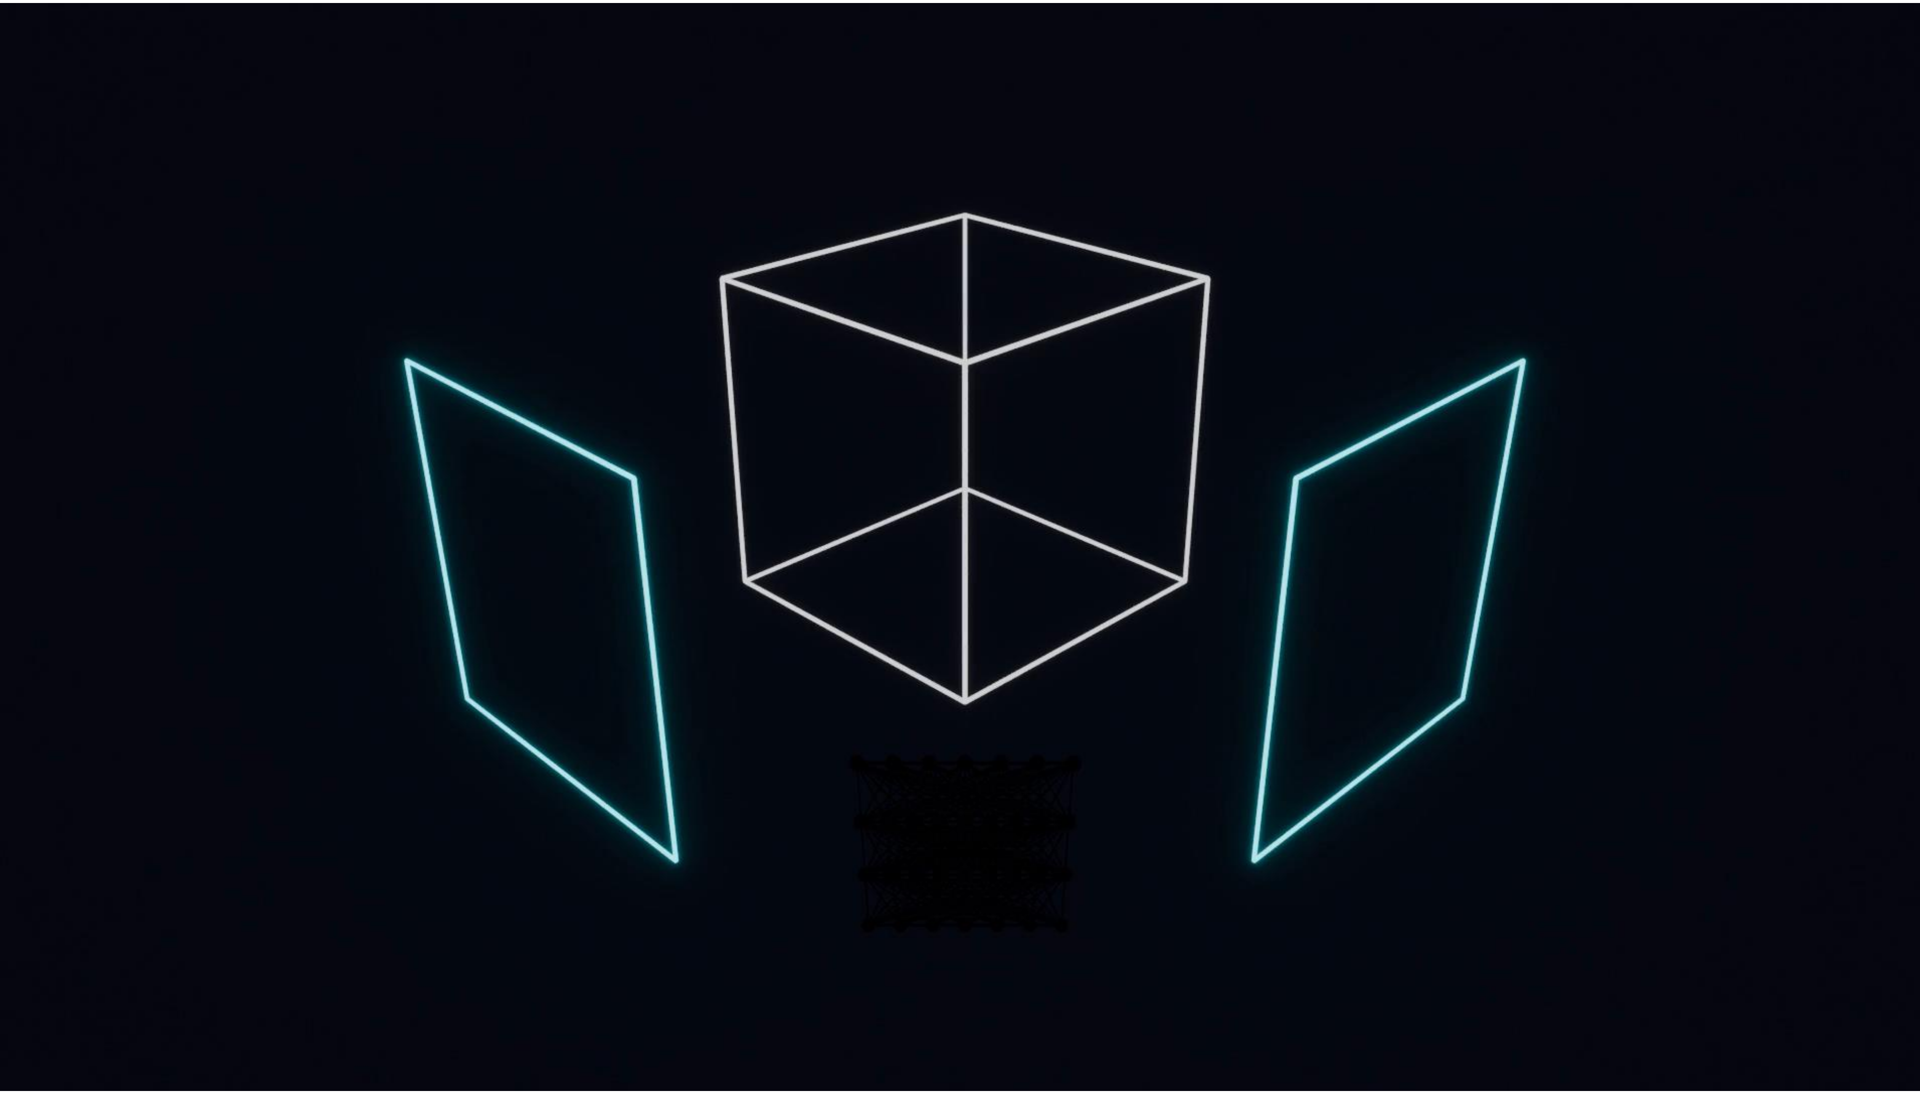

doi.org/10.1038/s41467-026-70918-x

Deep learning reconstruction in PLATON

- Using a Transformer encoder (inspired by *Google's BigBird*):
  - — **Tokenisation**: each 2D detected photon represents a token.
  - — **Sparse attention**: linear *O(n)* instead of quadratic *O(n^2)* complexity!
  - — **Output**: for each photon, the model outputs whether it's a dark count or a signal photon + its 3D scintillator origin.
  - — **Training MC sample validated using prototype data!**
- First trained on a sample with 1M simulated neutrino interactions:
  - — 10 x 10 x 10 cm^3 scintillator.
  - — Two orthogonal views instrumented.
  - — Optimised optics.

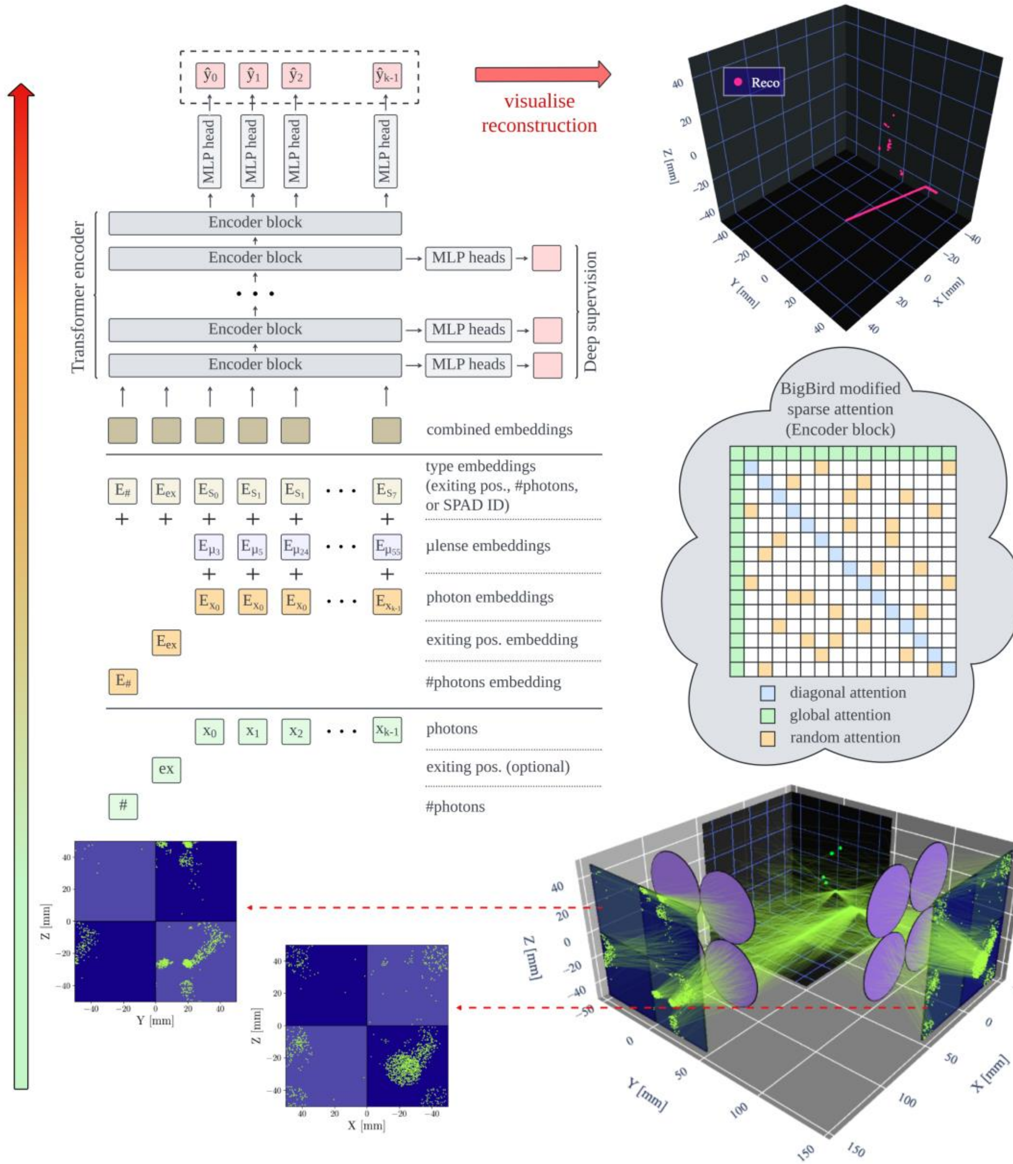

Training details

- 60-10-30% train/val/test splits.
- Embedding size of 384.
- Sequence length: maximum of 2048 photons per event.
  - For events with less than 2048 photons, random sampling 90% of the photons as a generalisation metric.
- 600 epochs with MSE loss, followed by 100 fine-tuning epochs with Huber loss.
  - Binary cross-entropy for dark-count/signal output.
  - Reconstruction loss only computed on signal pixels.
- Adam optimiser (40 epochs warmup + cosine decay scheduling).
- Effective batch size of 1024 events.
- 12 transformer-encoder layers with deep supervision.
- ~37M learnable parameters.
- Block size of 64 tokens.
- 4 random blocks per layer.
- 1 extra block of learnable latents.
- 32 GH200 GPUs (DDP, 8 nodes with 4 GPUs each, ~3 days of training).

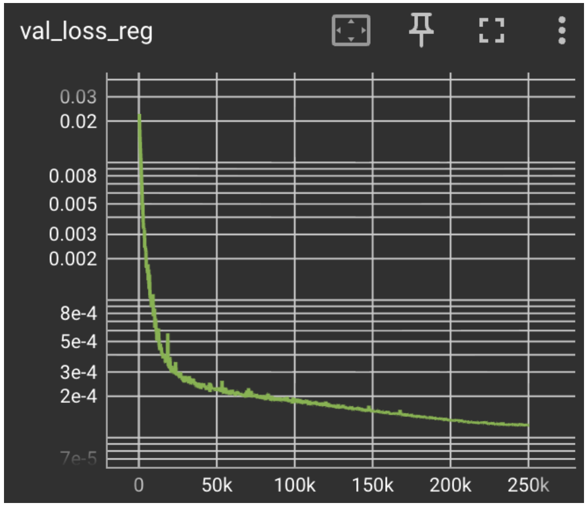

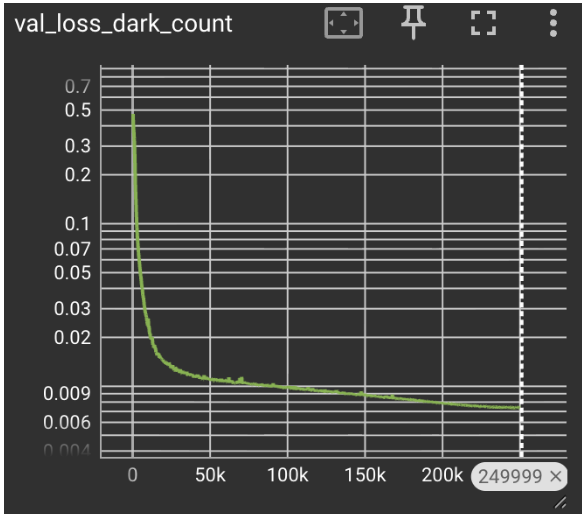

Inference: repeated subsampling and robust aggregation

1. Run the network multiple times per event.
   - Each pass uses a different random subset of up to 2048 detected photons.
2. Ensure enough predictions per photon
   - The number of passes is chosen automatically to reach approximately 99.9% photon coverage and at least about 5 predictions per photon.
3. Estimate uncertainty from repeated predictions
   - For each detected photon, collect all predicted 3D origins. The spread of these predictions gives an uncertainty estimate.
4. Filter and combine:
   - Reject photons with unusually large uncertainty using an event-adaptive threshold. For the remaining predictions, take the median 3D origin.

Event displays

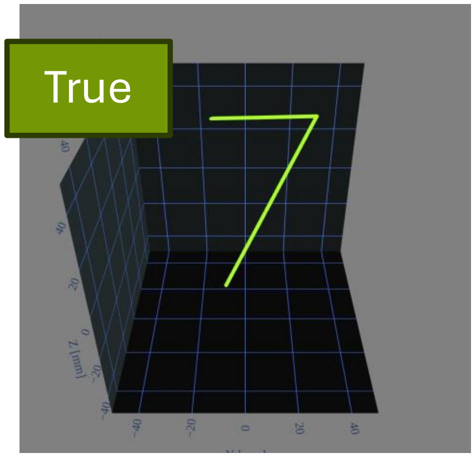

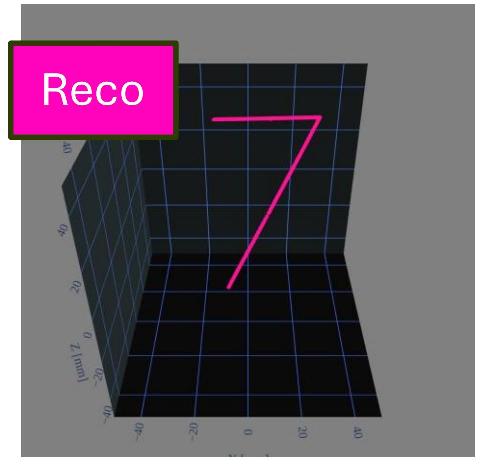

<200 μm 3D spatial resolution!

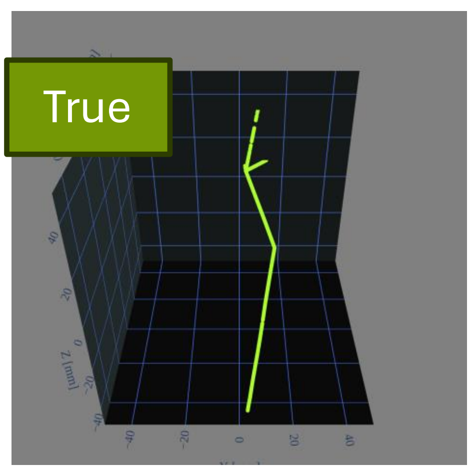

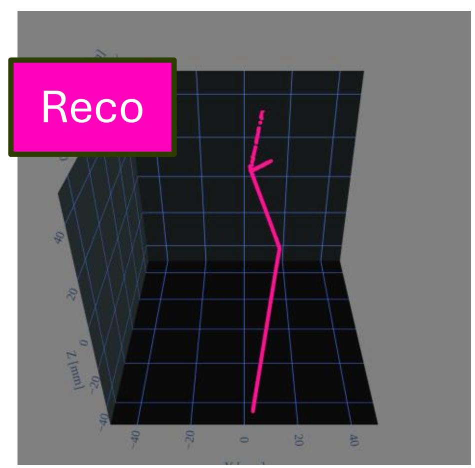

99.9% accuracy in discriminating between signal photons and dark counts!

doi.org/10.1038/s41467-026-70918-x

10 x 10 x 10 cm^3 results

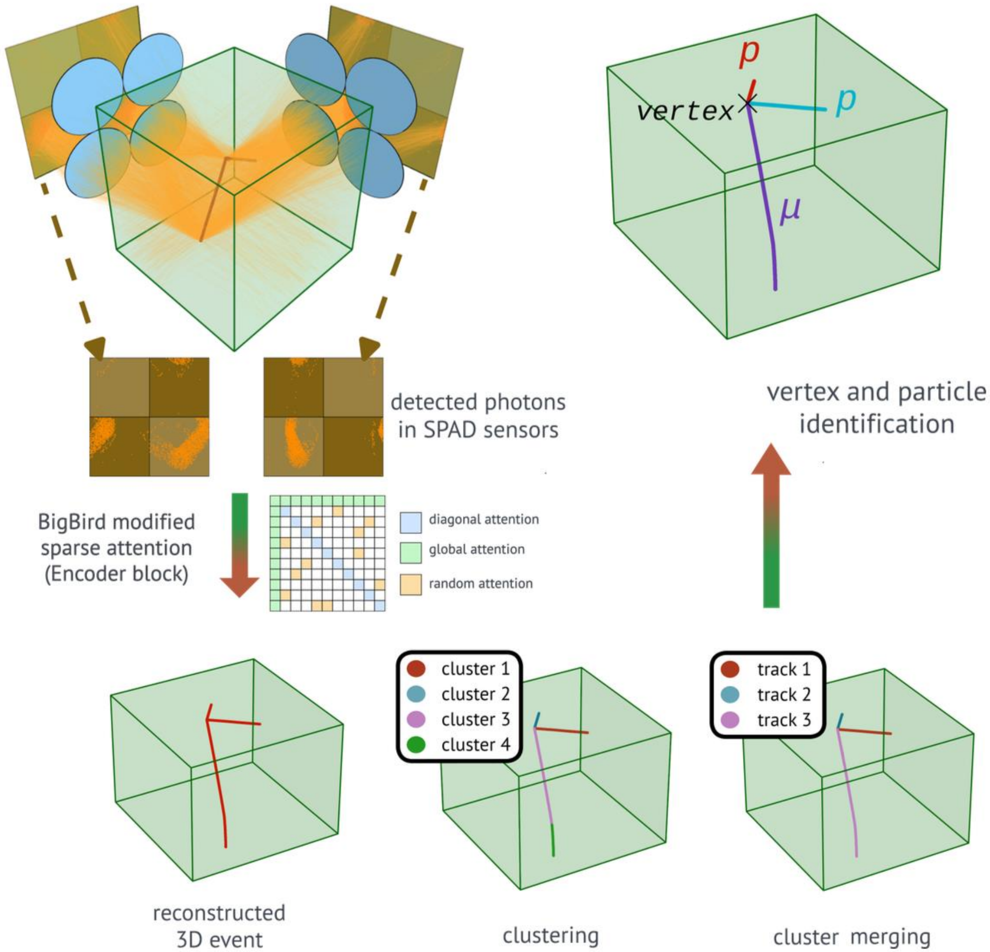

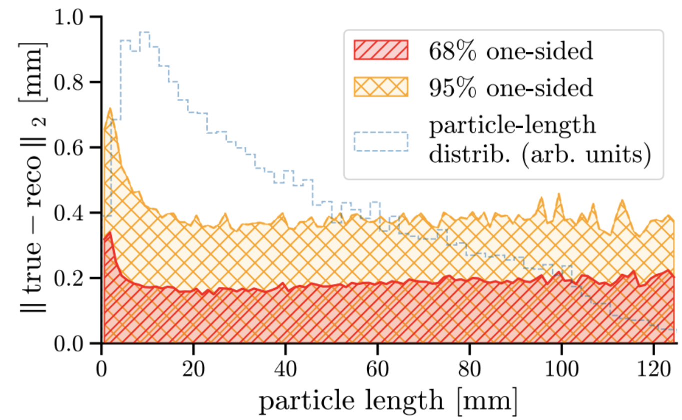

- Cluster the reconstructed point cloud into tracks, then perform physics analysis.

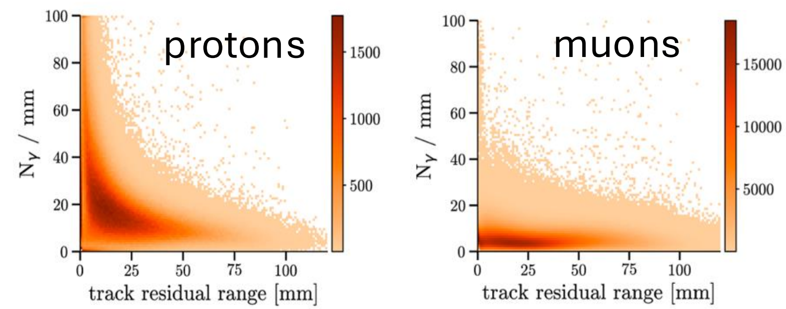

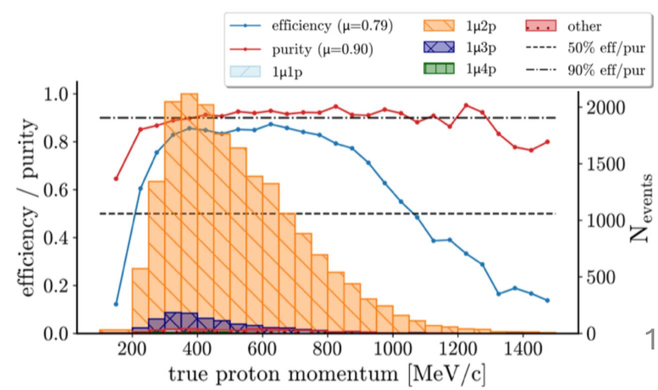

Scaling up: 1 x 1 x 1 m^3 scintillator

- At the tokenisation level, the representation remains unchanged:
  - Photon 2D position (relative to SPAD array).
  - Microlens centre (relative to SPAD array).
  - SPAD array centre.

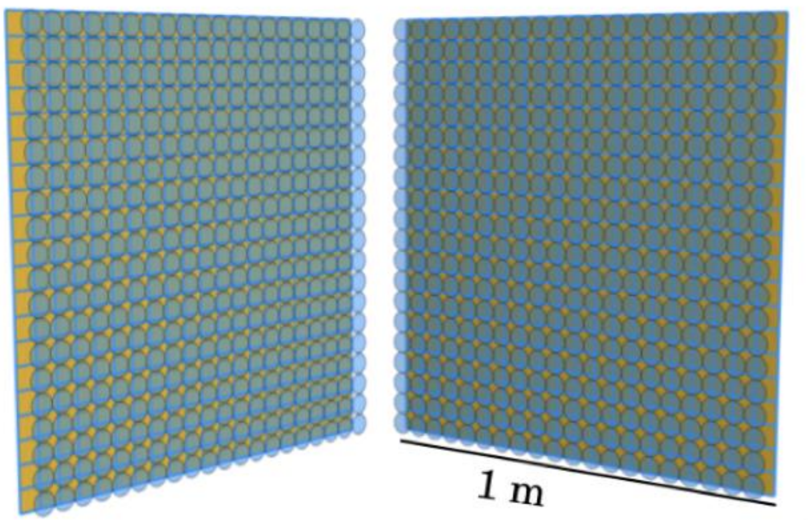

800 main lenses in total!

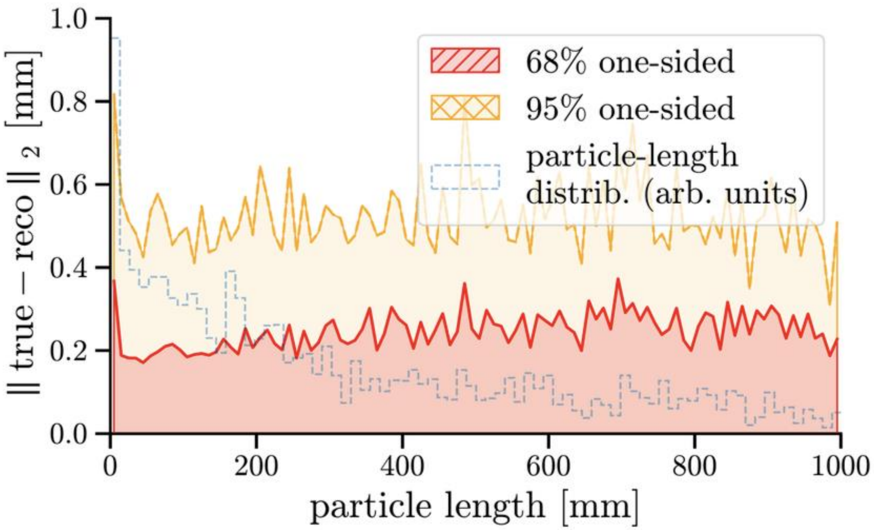

Transformer reconstruction achieves <300 μm 3D spatial resolution.

Example of electromagnetic shower reconstruction

Simulated EM shower in unsegmented 1 x 1 x 1 m^3 organic scintillator:

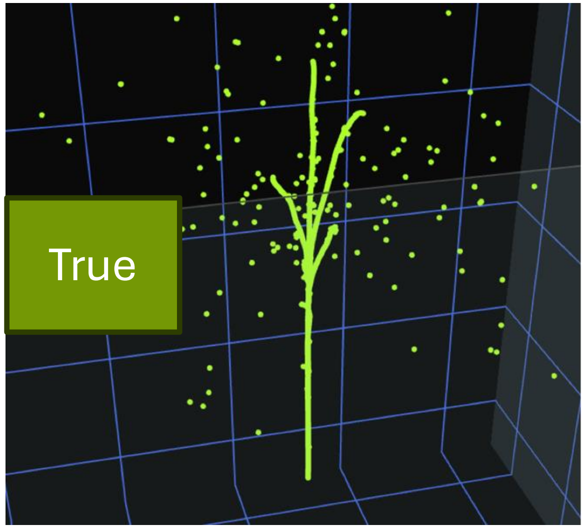

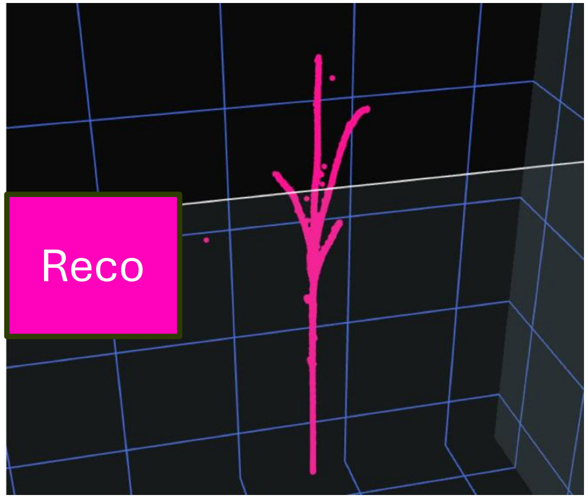

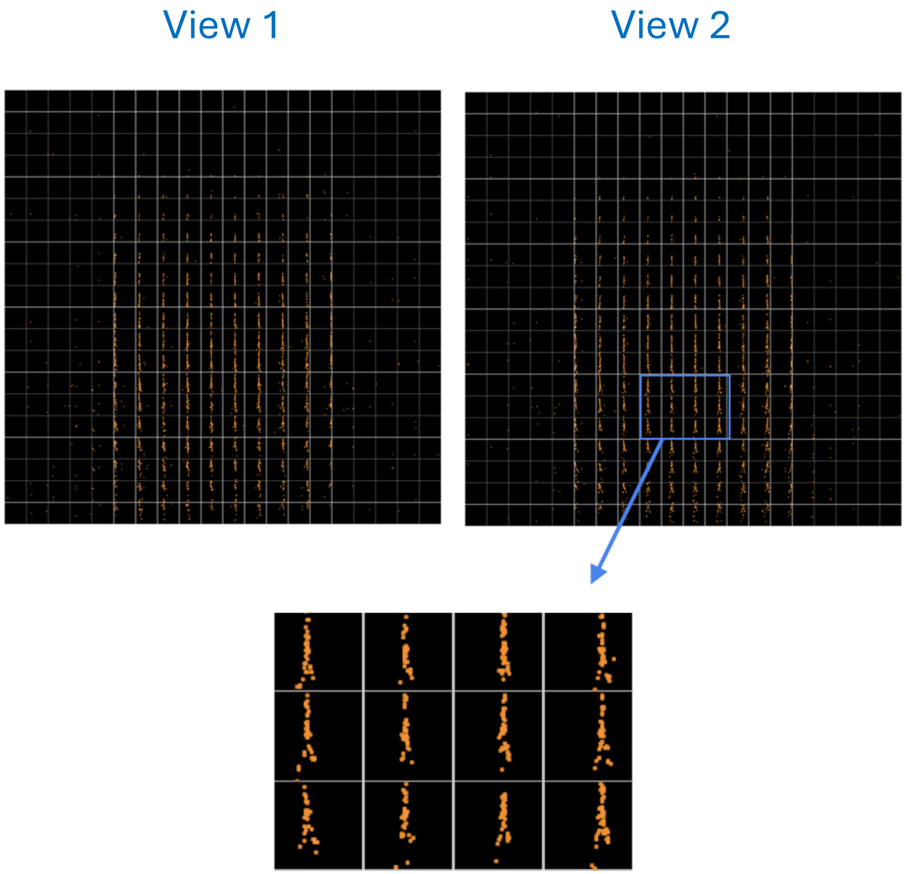

Generalisation test: non-neutrino image reconstruction

Simulated ETHZ logo in unsegmented 1 x 1 x 1 m^3 organic scintillator:

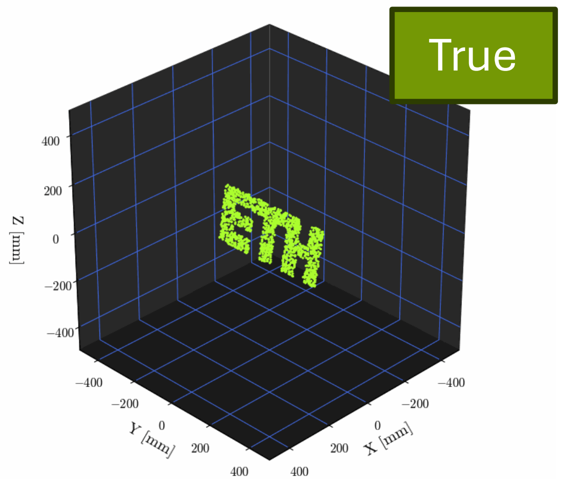

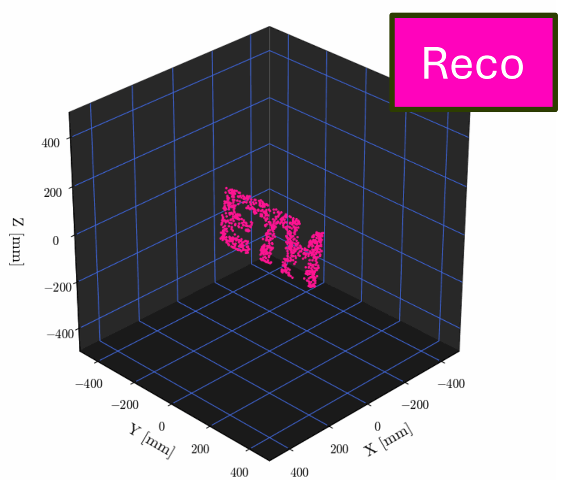

View 1

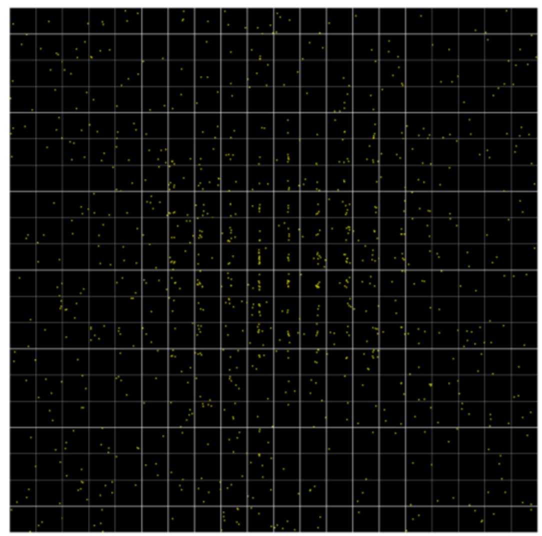

View 2

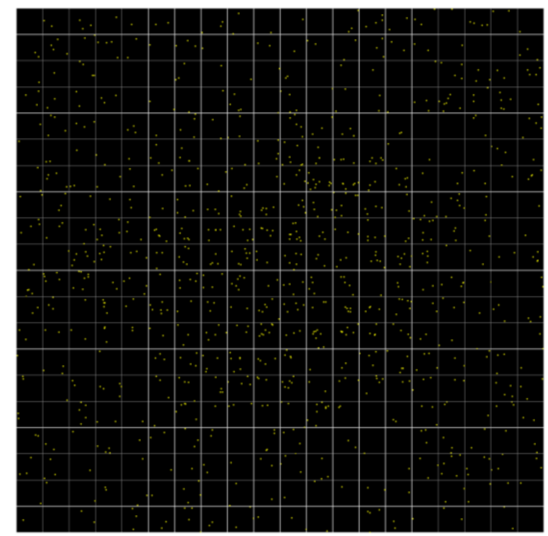

Model trained on neutrino interactions!

Seems to have learned the optics!

Summary

- Demonstrated 3D imaging of scintillation light with plenoptic cameras in starved-photon environment for particle detection.
- First ever made plenoptic camera instrumented with a SPAD array, with performance well reproduced in custom optical simulation.
- Deep-learning reconstruction resolves the optics.
- Studying the scalability of the PLATON concept to multi-meter scales.
- Currently working on fine-tuning the pre-trained model on different downstream tasks including PID and energy reconstruction.

DEEP LEARNING FOR PLENOPTIC NEUTRINO EVENT RECONSTRUCTION

Dr. Saúl Alonso-Monsalve

Neutrino Physics and Machine learning (NPML)

ETH Zürich

16 June 2026