

# **AI/ML Opportunities for Large Scale Liquid Scintillator Detector**

**Aobo Li**  
**Halicioğlu Data Science Institute**  
**Department of Physics**  
**UC San Diego**

**U. Mainz Liquid Scintillator Town Hall Meeting 10/13/2026**

# DOE Genesis Mission

Keywords from all funded HENP Genesis Mission Phase I Award



# Outline

## AI/ML directions relevant to Liquid Scintillator Detector

- SS/MS Separation in Liquid Scintillator: **The low-hanging Fruit**
- Foundation Model for Large-Scale Detector: **One model for all**
- An AI-Ready Datasets and Benchmarks: **Comparing Apple to Apple**
- Fast ML on FPGA: **The next-generation DAQ system**
- LLM and AI Agent: **Automating everything in close-loop**

# SS/MS Separation in Liquid Scintillator: **The low-hanging Fruit**

**TL;DR:** AI/ML can achieve Single-site vs. Multi-site classification in Monolithic Liquid Scintillator Detector

KamNet [[Phys. Rev. C 107, 014323](#)] shows that in KamLAND-Zen, neural networks can **separate  $\beta$  events from  $\gamma$  events**


- SNO+ has a likelihood-based method for SS vs. MS classification

This is a **new background rejection channel**, independent and multiplicative to coincidence tagging

- Adapted in KLZ  $0\nu\beta\beta$  results [[Phys. Rev. Lett. 130, 051801 \(2023\)](#); [Phys. Rev. Lett. 135, 262501 \(2025\)](#)]
- Combine with coincidence tagging method to further reduce long-lived spallation products

In our new work on DSNB search [[Accepted by ApJ](#)], we use KamNet to improve neutron tagging


- Application beyond  $0\nu\beta\beta$



*To harness this benefit, we need to create the best ML model architecture for this task*

# Foundation Model for Large-Scale Detector: **One model for all**


TL;DR: one model, many data, different task



# Foundation Model for Large-Scale Detector: **One model for all**


**TL;DR:** Foundation model for LArTPC and Water Cherenkov [\[https://arxiv.org/pdf/2609.00611\]](https://arxiv.org/pdf/2609.00611)

Credit: Kazuhiro Terao




Raw Detector Data

Learned Sensor-level Representations



Liquid Argon Time Projection Chamber



| True Ring Count | Predicted Ring Count |       |       |
|-----------------|----------------------|-------|-------|
|                 | 1R                   | 2R    | 3R    |
| 1R              | 96.0%                | 3.5%  | 0.5%  |
| 2R              | 8.1%                 | 90.1% | 1.8%  |
| 3R              | 0.2%                 | 2.6%  | 97.2% |

| True Number of Rings | fiTQun Reconstruction |       |           |
|----------------------|-----------------------|-------|-----------|
|                      | 1R                    | 2R    | $\geq 3R$ |
| True 1R              | 95.0%                 | 4.64% | 0.41%     |
| True 2R              | 27.8%                 | 66.7% | 5.56%     |
| True $\geq 3R$       | 7.04%                 | 25.5% | 67.5%     |

fiTQun on fully contained atmospheric (not applies to apples) [\[arXiv:1901.03230\]](https://arxiv.org/abs/1901.03230)

# Foundation Model for Large-Scale Detector: **One model for all**


**TL;DR:** Multimodal Foundation Model for different 0νββ detectors [Submitted to ICLR 2027]  
Genesis Mission Phase I Award: **lead PI Zepeng Li [U. Hawaii Manoa]**

Open, AI-Ready datasets from 4 0νββ experiments

Created 2 benchmarking metrics:

- **AUC<sub>w</sub>**: Energy-matched ROC AUC
- **I**: energy independence score

Tested different geometric ML models vs a unified transformer approach



| Model                                 | #Params<br>(K) | NEXT             |              | MJD                |                    | EXO-200          |              | SuperNEMO        |              |
|---------------------------------------|----------------|------------------|--------------|--------------------|--------------------|------------------|--------------|------------------|--------------|
|                                       |                | AUC <sub>w</sub> | <i>I</i>     | AUC <sub>w</sub>   | <i>I</i>           | AUC <sub>w</sub> | <i>I</i>     | AUC <sub>w</sub> | <i>I</i>     |
| <i>Native Representation Approach</i> |                |                  |              |                    |                    |                  |              |                  |              |
| Multi-view CNN                        | 808            | 0.977            | 0.971        | 0.970              | 0.760              | 0.863            | 0.752        | 0.586            | 0.910        |
| Static GINE                           | 480            | <b>0.995</b>     | 0.970        | 0.969 <sup>†</sup> | 0.797 <sup>†</sup> | <b>0.949</b>     | <b>0.801</b> | <b>0.590</b>     | 0.913        |
| BiGRU                                 | 734            | 0.979            | <b>0.971</b> | <b>0.975</b>       | <b>0.775</b>       | 0.929            | 0.789        | 0.579            | 0.921        |
| PointMamba-lite                       | 317            | 0.925            | 0.971        | 0.960              | 0.773              | 0.934            | 0.786        | 0.579            | <b>0.922</b> |
| <i>Unified Transformer Approach</i>   |                |                  |              |                    |                    |                  |              |                  |              |
| Entity + MLP                          | 111            | 0.975            | <b>0.975</b> | 0.929              | <b>0.812</b>       | 0.941            | 0.737        | 0.591            | 0.975        |
| Entity + Fourier                      | 113–116        | 0.963            | 0.974        | 0.941              | 0.802              | 0.947            | 0.729        | 0.589            | 0.974        |
| Entity + RoPE                         | 106–113        | 0.497            | 0.512        | 0.793              | 0.749              | 0.946            | 0.715        | 0.551            | 0.865        |
| Region + MLP                          | 111–121        | <b>0.988</b>     | 0.973        | 0.973              | 0.798              | 0.970            | 0.801        | 0.591            | 0.976        |
| Region + Fourier                      | 113–125        | 0.979            | 0.973        | <b>0.981</b>       | 0.788              | <b>0.972</b>     | 0.790        | 0.590            | 0.977        |
| Region + RoPE                         | 107–116        | 0.512            | 0.483        | 0.916              | 0.768              | 0.971            | 0.789        | <b>0.602</b>     | 0.953        |
| Summary + MLP                         | 111            | 0.987            | 0.971        | 0.965              | 0.799              | 0.947            | 0.776        | 0.590            | 0.975        |
| Summary + Fourier                     | 113–116        | 0.983            | 0.973        | 0.969              | 0.782              | 0.960            | <b>0.802</b> | 0.589            | <b>0.977</b> |
| Summary + RoPE                        | 106–114        | 0.602            | 0.532        | 0.878              | 0.777              | 0.955            | 0.773        | 0.582            | 0.956        |

Monolithic Liquid Scintillator Detector is missing from this table due to lack of Open, AI-Ready Datasets

# An AI-Ready Datasets and Benchmarks: **Comparing Apple to Apple**

**TL;DR:** the Monolithic Liquid Scintillator detector community needs an Open, AI-Ready Datasets and Benchmarks


Why do we need this datasets?

- Allowing different collaborations/experiments to compare their AI model **Apple-to-Apple**
- Reserve us a seat in the increasing trend of **HENP foundation model**

A few key decisions to make:

- What data? Simulation is okay, real data better
- What input? Point Cloud/Graph vs. Spatiotemporal
- What task? SS vs. MS Classification, Energy Regression, others

Within the KamLAND-Zen collaboration, we have one such dataset, but it's not open to the public





*If any experimental collaboration is interested in creating this dataset, I'm happy to talk and provide insights.*

# Fast ML on FPGA: The next-generation DAQ system

**TL;DR:** deploying AI algorithms onto the DAQ system

[NeurIPS ML4PS 2024, <https://arxiv.org/abs/2410.02991>]

*Current goal: significantly pushing down the inference latency from 0.5s to sub-ms level*




# Fast ML on FPGA: The next-generation DAQ system

**TL;DR:** deploying AI algorithms onto the DAQ system



FastML is a large community

- LHC, EIC and DUNE has significant investment in this direction
  - Genesis Mission Phase I Award: [Lead PI Wei Xie \[Purdue\]](#)
  - Genesis Mission Phase I Award: [Lead PI Kate Scholberg \[Duke\]](#)
  - Genesis Mission Phase I Award: [Lead PI Keith Ulmer \[U. Colorado\]](#)
- LEGEND has started in this direction recently
  - Genesis Mission Phase I Award: [Lead PI Richard Saldanha \[PNNL\]](#)
  - Deploying ML algorithms onto FPGA for real-time data cleaning
- KamLAND-Zen has been funded by NSF to investigate in this direction
  - BU, VirginiaTech, UPenn, UCSD, U. Hawaii
  - ML algorithms for real-time neutron tagging, position/ energy reconstruction, and more

FastML Foundation  
[<https://fastmachinelearning.org/>]



FastML For Science Conference 2026 at UCSD  
[<https://indico.cern.ch/event/1654479/overview>]




NSF A3D3 Institute  
[<https://a3d3.ai/>]

# LLM and AI Agent: Automating everything in Close-Loop

TL;DR: Large Language Model can orchestrate different tools and do research on its own

## LLM Agents for Auto-ML



**AI Scientist: End-to-end Auto ML Research**

**Latest achievement: \$15 per workshop paper on top AI conferences**

11

nature

[Explore content](#) [About the journal](#) [Publish with us](#)

[nature](#) > [articles](#) > article

Article | [Open access](#) | Published: 25 March 2026

### Towards end-to-end automation of AI research

[Chris Lu](#), [Cong Lu](#), [Robert Tjarko Lange](#), [Yutaro Yamada](#) [✉](#), [Shengran Hu](#), [Jakob Foerster](#), [David Ha](#) [✉](#) & [Jeff Clune](#) [✉](#)

[Nature](#) **651**, 914–919 (2026) | [Cite this article](#)

**307k** Accesses | **39** Citations | **823** Altmetric | [Metrics](#)

#### Abstract




The automation of science is a long-standing ambition in artificial intelligence (AI) research<sup>1,2</sup>. Although the community has made substantial progress in automating individual components of the scientific process, a system that autonomously navigates the entire research life cycle—from conception to publication—has remained out of reach. Here we present a pipeline for automating the entire scientific process end to end. We present The AI Scientist, which creates research ideas, writes code, runs experiments, plots and analyses data, writes the entire scientific manuscript, and performs its own peer review. Its ideas, execution and presentation are of sufficient quality that the manuscript generated by this AI system passed the first round of peer review for a workshop of a top-tier machine learning conference. The workshop had an acceptance rate of 70%. Our system leverages modern foundation models<sup>3,4,5</sup> within a complex agentic system. We evaluate The AI Scientist in two settings: a focused mode using human-provided code templates as an initial scaffold for conducting research on a specific topic and a template-free, open-ended mode that leverages agentic search for wider scientific exploration<sup>6,7</sup>. Both settings produce diverse ideas and automatically test, report on and evaluate them. This achievement demonstrates the growing capacity of AI for making scientific contributions and signifies a potential paradigm shift in how research is conducted. As with any impactful new technology, there could be important risks, including taxing overwhelmed review systems and adding noise to the scientific literature. However, if developed responsibly, such autonomous systems could greatly accelerate scientific discovery.

# LLM and AI Agent: Automating everything in Close-Loop

**TL;DR:** SIDERIUS is an resource-aware AI agent for scientific tasks

[Github: [yuema137/SIDERIUS](https://github.com/yuema137/SIDERIUS) ; Paper: <https://zenodo.org/records/23071121>]

Genesis Mission Phase I Award: [Lead PI Aobo Li \[UCSD\]](#)



*SIDERIUS is completely open source and works on any types of input data & Task, people without ML experience can use it to build AI model!*

— Claude CLI (Run 1) — FW, base (Run 3) — Orch., base (Run 5)  
— Codex CLI (Run 2) — FW, w/ DA<sup>+</sup> (Run 4) — Orch., w/ DA<sup>+</sup> (Run 6)

# Summary

## AI/ML directions relevant to Liquid Scintillator Detector

- SS/MS Separation in Liquid Scintillator:

**The low-hanging Fruit:** A demonstrated pathway about how AI can help

- Foundation Model for Large-Scale Detector:

**One model for all:** there is an increasing interest in HENP about Multimodal Foundation Model

- An AI-Ready Datasets and Benchmarks:

**Compare Apple-to-Apple:** Identified the best model for SS/MS separation

Secure our seats in the HENP foundation model trend

- Fast ML on FPGA:

**The next-generation DAQ system:** established research direction we could benefit from

- LLM and AI Agent:

**Automating Everything in Closed Loop:** SIDERIUS can design algorithm by itself

LLM and Agent is a tool to “glue everything together”

**Thanks to Funding  
Agency that Support  
Our Group!**



**Personal Website:**

[https://aobol.github.io/  
AoboLi/](https://aobol.github.io/AoboLi/)

**Email Address:**

[aol002@ucsd.edu](mailto:aol002@ucsd.edu)