Hi all,

Here's what you folks might find interesting today.

Anomaly detection and weak supervision

- **Weakly supervised machine learning for model-agnostic searches of new phenomena in the gamma-ray sky**

This is the clearest match in today's list to **model-agnostic anomaly detection**. The paper studies a background-versus-mixture setup, where the classifier is trained only on samples with different signal admixtures rather than explicit signal labels, and then tests that idea on pulsar/AGN separation, dark-matter subhalos, and axion-photon spectral distortions. What makes it worth a look is that the authors are really probing when weak supervision remains useful as the signal fraction and sample composition vary, which is exactly the practical question that comes up in weakly supervised searches. This might be interesting for you, Aritra, because even though the application is gamma-ray data rather than the LHC, the methodological overlap with **weakly supervised, model-agnostic new-physics searches** is direct. Jingjing, you might also care a bit about it from the broader inference angle, but this one is much more squarely in Aritra's lane. No author boost here, but topically it is still the strongest paper today.

Graph ML for detector and event reconstruction

- **Machine learning pipeline for identifying tracks of muons and hadrons at GRAPES-3 muon telescope**

This paper builds a full ML chain for particle identification in a cosmic-ray detector, moving from single-particle classifiers to event-level graph neural networks where detector hits are nodes and node labels correspond to muon versus hadron hits. The useful part is not just that they use a GNN, but that they formulate the multiparticle classification problem at the hit level with edge-convolution layers, then follow it with a regression model to estimate total particle and muon multiplicities. This might be interesting for you, Aritra, because it is another example of **graph-based architectures for sparse detector data**, even if it is not directly about jets or LHC anomaly searches. It is less aligned with your highest-priority topics than hypergraphs or jet-substructure ML, so I would treat it as a secondary pick rather than must-read. Still, if you are tracking how graph methods are being operationalized in HEP-style reconstruction pipelines, there is enough concrete architecture detail here to be useful.

- **The Deep Learning Cosmic Ray Energy Reconstruction Pipeline for the GRAPES-3 Experiment**

This one is another detector-ML paper, but the angle is energy reconstruction with a modular GNN rather than classification. The potentially useful detail is that the authors explicitly study the latent space learned by the network and argue that rescaling the latent metric improves response resolution, so there is at least some attempt to understand what the model is encoding rather than presenting a black-box benchmark. This might be interesting for you, Aritra, because it touches the broader question of how graph architectures represent event geometry and detector response, which is relevant to modern collider ML even outside the jet context. It is not a close match to your top-priority topics, and it is not especially targeted to Jingjing's inference interests either. I would file it as a "skim if you have time" paper rather than a priority read.

Agentic AI, automation, and analysis infrastructure

- **Data Preservation in High Energy Physics: Global Report 2026**

This is not an ML methods paper, but it is probably the most relevant infrastructure piece for **AI-enabled workflows** in today's batch. The report emphasizes a shift toward automation in data preservation, including AI/ML for metadata extraction, curation, and workflow optimization, and it frames that in the context of long-term analysis sustainability rather than one-off tooling. This might be interesting for you, Jingjing, because one of your priority areas is **agentic AI and workflow automation for particle-physics analysis tasks**, and this report gives a broad view of where the community is actually moving operationally. It is not about MC generation or code-writing agents specifically, so it is not a perfect match, but it does speak directly to the institutional and technical ecosystem those tools would have to live in. I would not put it in the must-read bucket unless you are actively thinking about automation infrastructure, but it is more relevant than most of today's non-ML phenomenology papers.

ML for reconstruction in rare-event and beyond-collider experiments

- **First Demonstration of a Hybrid Cherenkov and Scintillation Detector in a Proof-of-Principle Axion Search at a Beam Dump**

The ML content here is fairly focused but concrete: the thesis develops machine-learning-based position reconstruction at about 5 cm resolution and energy reconstruction at about 10 percent resolution for MeV-scale events in a hybrid optical detector. The more interesting technical point is that these reconstructions are embedded in a broader likelihood-based background rejection strategy using timing, wavelength sensitivity, directionality, pulse shape, and topology to isolate ALP-like events. This might be interesting for you, Jingjing, because it sits near your broader interest in **uncertainty-aware reconstruction and analysis workflows**, though it is not an unfolding or SBI paper. Aritra, you might also find it mildly relevant from the perspective of ML-assisted event characterization in nonstandard detector systems, but it is not close to your main collider-ML themes. Overall, this is a niche but respectable cross-disciplinary pick if you want to see how ML is being integrated into low-energy rare-event searches.

Physics-informed neural methods in fundamental theory

- **Neural-Spectral Discovery of Rotating Black Holes Beyond General Relativity**

This paper is outside particle-physics phenomenology, but it is one of the more technically serious ML papers in the list. The authors combine physics-informed neural networks with a pseudo-spectral refinement step to produce continuous neural-field solutions for rotating black holes in higher-curvature gravity, with residuals checked to high precision. The reason I am flagging it at all is that the hybrid PINN plus spectral-refinement workflow is more rigorous than many "neural PDE solver" papers, and the idea of certifying the learned solution rather than just fitting it is the main novelty. This might be interesting for you, Jingjing, if you are broadly tracking ML methods for scientific inference and simulation, but it is definitely peripheral to your core HEP priorities. I would only read it if you want methodological inspiration rather than direct application.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 31 papers hep-ex: 13 papers hep-th: 36 papers