Hi all,

Here's what you folks might find interesting today.

TOP PICKS

Jet tagging with novel ML architectures / ML exploiting jet substructure or deeper jet physics

ParticleTransformer is all you need for reconstructing hadronic tau leptons

This is probably the clearest ML-for-reconstruction paper in today's list. The authors take *ParticleTransformer* and push it beyond tagging into a full hadronic tau reconstruction stack: identification, decay-mode classification, charge reconstruction, and four-momentum regression, all in a realistic FCC-ee setup. What stands out is the comparison between task-specific models and a unified multitask model, with strong gains over conventional reconstruction-level observables and jet-charge estimators. For you, Aritra, this feels relevant because it is very much in the spirit of learned object reconstruction from particle-level inputs, with architectural lessons that could transfer to jet tagging and substructure-heavy LHC problems. Even though it is taus rather than jets, the core message is that transformer-style particle representations can replace hand-built heuristics across several coupled reconstruction tasks. If you are thinking about how far modern set/particle architectures can be pushed in collider reconstruction, this is the one I would open first.

Anomaly detection at the LHC / Novel ML or statistics methods with possible HEP use

Stress testing of fast reconstruction pipelines using machine learning

This paper is not anomaly detection in the usual new-physics-search sense, but it does use an *unsupervised regime-mapping framework* to diagnose and repair failures in fast reconstruction pipelines. The key idea is to stress-test the common "local response" assumption in fast simulation and reconstruction, then explicitly expose cases where global event context matters and biases the reconstructed observables. Their HL-LHC example with $Z \to \ell\ell$ shows a concrete failure mode: peak instability, bias, and degraded resolution, followed by recovery once the unsupervised mapping is used. This might be interesting for you, Aritra, because it is a model-agnostic unsupervised method aimed at finding hidden structure in detector-response behavior rather than signal labels. It is also potentially useful for Jing, since robust simulation and reconstruction diagnostics matter directly for inference pipelines that rely on approximate simulators. I would read this as a methods paper about where ML can audit the assumptions built into fast HEP workflows.

Agentic AI and workflow automation for particle physics analysis tasks / General ML methods

ToolChain-CRC: Conformal Risk Control for Agentic AI Under Retrieval and Tool-Use Drift

This is the most directly relevant agentic-AI paper in the batch. The novelty is that it does not just calibrate the final answer of an AI agent, but the *entire trajectory* of retrieval steps, tool calls, intermediate observations, and final output, with an accept-or-intervene rule and an anytime alarm for early stopping. That matters because workflow automation in physics will likely fail through bad retrieval, wrong tool outputs, or unsupported intermediate steps long before the final text looks suspicious. This might be interesting for you, Jing, especially if you are thinking about agentic systems for code writing, sample-generation orchestration, or analysis automation where auditability and failure control are essential. It is not particle-physics-specific, but it is exactly the

kind of reliability layer that would make agentic analysis tools more usable in practice. I would file this under enabling infrastructure for safe automation rather than direct physics ML.

Novel ML or statistics methods with possible HEP use

Bayesian Nonparametric Detection of Anomalies in Multivariate Functional Data

This is a general anomaly-detection methods paper, but it has enough substance to be worth a look. The method models functional data with an infinite mixture of multi-output Gaussian processes, using wavelet-based mean functions, Besov priors for smooth sparse structure, and an intrinsic coregionalization model for cross-function dependence. The practical hook is that anomalies are identified as observations assigned to small mixture components, so the framework does not need the anomaly class to be specified in advance. This might be interesting for you, Aritra, because the setup is naturally compatible with high-dimensional structured observables where "normal" behavior is multi-modal and rare departures are the target. It may also be useful for Jing as a statistical template for weakly labeled or semi-supervised settings where one wants uncertainty-aware anomaly scoring. Not directly collider-specific, but more methodologically serious than a generic anomaly paper.

Novel ML or statistics methods with possible HEP use

Optimal score function estimation via derivatives constraints

This one is fairly theoretical, but it touches a toolset that keeps showing up in modern simulation and inference. The authors study score-function estimation through empirical risk minimization and argue that constraining the hypothesis class to a Sobolev ball is enough to avoid overfitting while achieving minimax rates; they then connect this to score-based generative modeling. This might be interesting for you, Jing, because score estimation is adjacent to simulation-based inference, likelihood-free methods, and generative surrogates used in high-dimensional analyses. The appeal here is not an immediate collider application, but a cleaner statistical understanding of when score-based methods are actually estimable and well-behaved. If you are tracking foundational work that could eventually sharpen SBI or generative simulation tools, this is a reasonable candidate. I would treat it as a theory paper to bookmark rather than an urgent read.

Quantum ML / quantum architectures with possible HEP relevance

Exponentially many initializations to avoid barren plateaus

Among the quantum-ML-adjacent papers, this is the strongest fit. The paper argues that barren plateau avoidance is not just about finding *an* initialization, but about navigating exponentially many inequivalent initialization families, each potentially steering training into different "trainable pockets." For you, Aritra, this could be relevant if you are thinking about Lorentz-equivariant or symmetry-aware quantum circuit constructions, since initialization pathologies are one of the main practical bottlenecks in trainable quantum models. It is not HEP-specific and does not mention equivariance, but it gives a more nuanced operator-level diagnostic for trainability than the usual folklore. The fact that different first-moment-distinct initializations can land in different minima is also a useful cautionary point for benchmarking quantum architectures. I would read this as a conceptual paper about optimization landscape structure in variational quantum models.

BY CATEGORY

Anomaly detection at the LHC / Novel ML or statistics methods with possible HEP use

Sequential Kernel-based Conditional Independence Testing via Adaptive Betting

FOSC-X: An Extended Framework for Optimal Local Cuts and Non-Horizontal Cluster Selection from Clustering Hierarchies

Agentic AI and workflow automation for particle physics analysis tasks

The Coherence Principle: A Falsifiable Prior for Model Selection from the Grammar of Theories

Jet tagging with novel ML architectures / ML exploiting jet substructure or deeper jet physics

High-$p_{\rm T}$ physics and jet production

Quantum ML / quantum architectures with possible HEP relevance

Quantum simulation of neutrino oscillations with bosonic encoding

Quantum Annealing Enhanced Reinforcement Learning for Accurate Remaining Useful Lifetime Prediction

General HEP, lower priority today

Electromagnetic Shower Reconstruction and Identification in FASER's Emulsion Detector for LHC Forward Neutrino Measurements

QCD studies and precision physics at the LHeC

The Z'-boson of the $B-L$ Supersymmetric Standard Model and its Large Hadron Collider Searches

Probing Long-Lived Particle Production in Muon Decays at the SNS with a Highly Capable Hydrocarbon Detector

Standard Candles for Supernova Neutrino Detection at DUNE

Until the next announcement!

P.S. In the most recent submission window: hep-ph: 30 papers hep-ex: 12 papers stat.ML: 31 papers