Hi all,

Here's what you folks might find interesting today.

Top picks

Researcher-priority topics

[Better Queries, Cheaper Attention: Adapting Transformers for Efficient Sparse Reconstruction]

This one is probably the clearest ML-for-HEP hit in today's list. The paper replaces fixed decoder queries with *input-conditioned dynamic queries* and combines that with *geometry-defined local strided cross-attention*, so the transformer only attends to physically plausible hit-track associations. That is a concrete architectural change, not just an incremental benchmark, and it matters because HL-LHC tracking is exactly the kind of sparse, high-multiplicity setting where standard attention becomes the bottleneck. They report both better reconstruction quality and much better efficiency: higher track efficiency, lower fake rate, about 50% lower inference latency, and more than an order-of-magnitude lower peak memory. This might be interesting for you, Aritra, as an example of detector-aware transformer design that could inspire similarly structured attention schemes for jet or event-level reconstruction problems. It is not about jets or anomaly detection directly, but it is one of the stronger examples here of novel ML architecture tailored to collider geometry.

[Model Validation of Agentic AI Systems: A POMDP-Based Framework for Belief-State, Forecast, and Policy Validation]

This is not particle-physics-specific, but it lines up well with the agentic AI and workflow automation angle. The main contribution is a validation framework for autonomous agents that decomposes performance into information acquisition, belief updates, forecasts, actions, and utility, all cast in a POMDP language. That is useful because a lot of current "agent" work is hard to evaluate rigorously beyond end-task success, and this paper gives a more structured way to audit where an agent is actually failing. For particle-physics workflows, that could translate naturally to automated analysis chains, code-writing agents, or MC-production orchestration where one wants to validate not just outputs but intermediate decision quality. This might be interesting for you, Jing, especially if you are thinking about agentic systems for analysis automation rather than just LLM wrappers. It is a methods paper rather than an application paper, but the framework feels transferable.

General methods with possible HEP relevance

[Nested Sampling: A Critical and Comprehensive Theoretical Guide]

This is a theory-and-practice guide to nested sampling rather than a new algorithm, but it is still relevant because nested sampling keeps showing up in inference-heavy parts of physics. The useful part is that it focuses on the actual hard step, namely sampling from likelihood-constrained priors, and tries to separate the clean derivation from the approximations people often gloss over. For simulation-based inference and broader likelihood-free workflows, that kind of clarity can be valuable when comparing evidence estimation or posterior exploration strategies. This might be interesting for you, Jing, as background reading if you are thinking about robust inference pipelines in high-dimensional settings. It is not SBI in the neural ratio estimation sense, so I would not rank it as a core match, but it is still adjacent enough to be worth a look. More broadly, it is one of the better "methods literacy" papers in today's set.

[Finsler Geometry, Graph Neural Networks, and You]

This is a mathematically flavored GNN paper that may be worth a skim for architecture ideas. The key point is that standard Laplacian-based GNNs are tied to isotropic diffusion, while the authors build graph layers from discrete approximations to the *Finsler Laplacian*, which can encode anisotropic geometry. That could matter in HEP settings where relational structure is directional or anisotropic, for example in detector geometry, shower development, or possibly jet constituent relations where Euclidean isotropy is too restrictive. This might be interesting for you, Aritra, especially given your interest in nonstandard graph constructions and architectures beyond vanilla message passing. It is not a hypergraph paper, but it does push geometric inductive bias in a direction that could be useful for collider applications. The novelty here is less about benchmark wins and more about expanding the operator class that GNN layers can represent.

FoundCause: Causal Discovery with Latent Confounders from Observational Data

This paper proposes an amortized causal discovery model that predicts causal graphs in one forward pass, including explicit handling of latent confounders. Architecturally, it mixes a permutation-invariant transformer, statistics-conditioned attention using classical asymmetry features, a factorized edge/direction decoder, and a dedicated confounder module with latent tokens. That combination is fairly novel, and the latent-confounder modeling is probably the most notable part. For HEP, I do not see an immediate direct application, but the general idea of amortizing expensive structure-learning over synthetic training distributions could be useful in simulation-rich scientific settings. This might be interesting for you, Jing, if you are broadly tracking ML methods that turn repeated inference or discovery tasks into learned surrogates. It is more speculative for particle physics than the top two picks, but methodologically it is one of the stronger ML papers here.

Fast Nonparametric Conditional Independence Testing via Two-Stage Regression

BLITZ is a fast conditional independence test designed for the repeated-query regime of constraint-based causal discovery, with a broad-to-local residualization strategy: first remove smooth dependence with low-order polynomials, then clean up nonlinear leftovers with shallow tree regressions. The practical claim is appealing: better calibration than other fast nonparametric tests while staying fast enough for thousands of CI queries. That is not directly particle physics, but scalable, well-calibrated independence testing can matter in scientific discovery pipelines and in diagnostics for learned simulators or inference systems. This might be interesting for you, Jing, particularly from the perspective of statistically reliable building blocks for automated analysis workflows. I would treat it more as a toolkit paper than a must-read, but the emphasis on calibration rather than just speed is a real plus. If you are thinking about structure learning or robust diagnostics, this is probably the one to keep.

Conformal Prediction Intervals with Tail-Specific Guarantees

This paper extends split conformal prediction to provide separately calibrated lower-tail and upper-tail guarantees, then intersects them into a two-sided interval. The main reason that matters is that many scientific problems care asymmetrically about one tail, and standard marginal coverage can hide poor directional calibration. In HEP contexts, that could be relevant anywhere uncertainty bands are used operationally and under-coverage in one tail is more costly than the other. This might be interesting for you, Jing, as a general uncertainty-quantification tool that could complement inference pipelines or surrogate models. It is not particle-physics-specific, but it is the kind of statistically careful method that can become useful once translated into analysis practice. I would file it under "good to know" rather than immediate priority.

By category

hep-ph

Search for a Time-Dependent Z' Resonance in the Dimuon Channel

PineAPPLv1: fast and flexible theory predictions for present and future colliders

Quantum decoherence of hyperon spin correlations in QCD hadronization

Quantum Resources and Wigner Symmetry in Nucleon-Nucleon Scattering from Effective Field Theory

Scaling of the Surface Free Energy as a Probe of the QCD Critical Region

Short-Range Correlations Between Partons in a Proton

Perturbative QCD as a quantitative tool in the years 1976-2000

Fourier-Preconditioned Path Deformations for Multi-Field Vacuum Tunnelling

When Renormalisation Remembers: UV/IR Mixing as an Entanglement Bridge

Supersymmetric geometry in non-supersymmetric effective field theory

hep-ex

Quantum decoherence of hyperon spin correlations in QCD hadronization

Optimized filtering for pulse-shape based pile-up rejection applied to 0nu beta beta search with 100Mo

Final Report on the Measurement of the Positive Muon Anomalous Magnetic Moment at Fermilab to 127 ppb

Improved limits on a new Z' in B-L scenarios with the NA64 experiment at CERN

Are neutrinos Majorana? Fixed-target and high-energy astrophysical searches decide

Direct Measurement of the 212Pb and 214Pb beta Decay Branching Ratios with the XENONnT Experiment

Semi-analytical results for e+e- to J/psi + X_non ccbar up to O(alpha_s v^2) at B factories

Observation of an Altered a0(980) Line shape in D+ to pi+ eta eta

stat.ML

Geometrical fairness in graph neural networks

Uncertainty Quantification of Engineering Structures by Polynomial Chaos Expansion and Multivariate Active Learning

Non-asymptotic Tail Bounds for the Kostlan-Shub-Smale Field: Tensor PCA and Spherical k-Spin Complexity

Differential Privacy of Gaussian Process Posterior Sampling

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 27 papers hep-ex: 9 papers stat.ML: 27 papers