Hello folks,

Here's what you might find interesting today. Fair warning: it's a slim batch, with most papers falling outside the core ML-for-HEP intersection. A few are worth flagging though.

Differentiable Physics Engines and Inference Infrastructure

MANGO: An Autodiff Neutrino Oscillation Engine for Differentiable Analysis Pipelines -- This might be interesting for you, Jingjing. The paper builds a fully differentiable neutrino oscillation probability engine where every quantity -- propagation geometry, detector depth, Earth-shell densities -- is differentiable with respect to all inputs. Because reverse-mode autodiff cost scales with output dimension, sensitivities for all 369 Earth-density parameters cost no more than the 6 standard oscillation parameters. The key novelty is that gradients flow continuously through the probability stage and onward through detector response, event weighting, binning, and likelihoods, enabling end-to-end gradient-based inference that traditional analytic formulas cannot achieve. This is directly relevant to your interest in **simulation-based inference** infrastructure, as it demonstrates how differentiable physics models can serve as the backbone for likelihood-free or gradient-assisted inference pipelines in neutrino physics.

Newtrinos.jl: A Julia Package for Global Analysis of Neutrino Data -- This is also relevant to you, Jingjing. The package provides a modular three-layer architecture separating physics models, experiment descriptions, and statistical inference, with all forward models automatically differentiable via Julia's AD ecosystem. It supports both Frequentist and Bayesian inference and defines both the likelihood and the data-generating process for each dataset, which is exactly the setup needed for **simulation-based inference** workflows. The composability of joint likelihoods across experiments and the automatic differentiability make this a practical tool for the kind of unbinned, high-dimensional analyses you care about, though it is neutrino-focused rather than LHC-focused.

Generative ML for Detector Simulation

Fine-tuned Normalizing Flows for ALICE Zero Degree Calorimeter Fast Simulation -- This one is for you too, Jingjing. The authors use normalizing flows with transfer learning and gradual unfreezing to build a generative surrogate for the ALICE ZDC neutron detector, replacing expensive Monte Carlo chains. What stands out is the introduction of physics-motivated evaluation metrics -- conditional weighted MAE, dispersion ratio, and Jaccard co-activation error -- that go beyond standard Wasserstein distance to capture conditional input-output dependencies that matter for downstream physics. This connects to your interests in **ML for simulation** and **workflow automation**, as it demonstrates a practical pipeline for accelerating detector response simulation with careful physics-aware validation.

Unfortunately, nothing in today's batch touches on anomaly detection, jet substructure, jet tagging, hypergraph networks, or Lorentz-equivariant architectures, so nothing to flag for you, Aritra. Better luck tomorrow.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 31 papers hep-ex: 11 papers hep-th: 32 papers