Hey there,

Here's what you folks might find interesting today.

ML for Event Generation and Simulation

Efficient Event Generation for High-Multiplicity LHC Processes: An End-to-End GPU Workflow with Normalizing Flows

This might be interesting for you, Jingjing. The paper presents the first end-to-end GPU-resident event-generation workflow that integrates **normalizing-flow** proposal densities with the parton-level generator Pepper, achieving up to two orders of magnitude speedup for high-multiplicity processes like $pp \to t\bar{t}+4j$ and $pp \to 5j$. The key novelty is the helicity-conditioned coupling flows trained with online updates and sample replay, exchanged directly in device memory between a Python control layer and Pepper. This directly addresses the **MC sample generation** bottleneck you care about, and the normalizing-flow architecture is closely related to the density estimation techniques used in **simulation-based inference**. The fact that they generate 10^9 unweighted events on four H100 GPUs in under a day is a concrete demonstration of ML-driven workflow automation for collider physics.

ML for Collider Physics

On-Detector Machine Learning for Beam-Induced Background Rejection at a 10 TeV Muon Collider

This might catch your eye, Aritra. The paper explores **on-detector ML** for beam-induced background rejection in the vertex detector of a future 10 TeV muon collider, using lightweight neural networks operating on pixel cluster shapes. Three classes of architectures are evaluated for FPGA/ASIC feasibility via high-level synthesis, achieving 88-90% data reduction at 99% signal efficiency. While this is not LHC physics per se, the problem of distinguishing collision products from overwhelming background using low-level detector information is conceptually adjacent to your interest in **model-agnostic searches** and **anomaly detection**. The hardware-constraint angle and the focus on real-time inference also distinguish this from typical offline ML approaches.

Di-Higgs Phenomenology

Higgs Boson plus Quark Pair Production at High Energies as di-Higgs Background

This is relevant for you, Jingjing, though it lacks an ML component. The paper provides the first dedicated study of $H+b\bar{b}+2j$ production in the gluon-fusion channel at high energies as a **background to VBF Higgs-pair production**, using **High Energy Jets (HEJ)** resummation with quark-pair effects included for the first time. The key finding is that fixed-order calculations overestimate the gluon-fusion contamination in VBF-like phase-space regions, which has direct implications for **HH sensitivity studies** at the LHC. If you are thinking about projected sensitivity for di-Higgs searches, this refines an important background estimate -- though you would likely want to pair it with an ML-based analysis strategy to fully leverage it.

Jet Substructure Physics

Unraveling QCD dynamics with heavy quark energy correlators

This might be interesting for you, Aritra, as it connects to your interest in **understanding the physics of jets**. The paper studies heavy-quark **energy-energy correlators (EECs)** that are explicitly sensitive to charm and bottom quark mass effects, showing how finite mass reshapes intra-jet radiation patterns and produces measurable deviations from massless expectations. The extension to nuclear matter, where medium-induced radiation competes with vacuum mass suppression, adds a novel dimension. While there is no ML component, the information-theoretic and statistical structure of EECs as jet substructure observables is the kind of physics that ML-based jet taggers ultimately exploit, so understanding these mass-dependent features could inform the design of more physically motivated architectures.

ML for Scientific Discovery

CDRL: Certification-Driven Reinforcement Learning for Neutrino Flavor Model Discovery

This is tangentially relevant for you, Jingjing, in the context of **agentic AI and workflow automation**. The paper introduces Certification-Driven RL, where symbolic reasoning tools generate structured feedback certificates that are converted into reusable constraints, guiding exploration away from invalid regions of a combinatorial hypothesis space exceeding 10^{26} models. The framework achieves up to 6.33x higher valid model rates than standard RL for neutrino flavor model discovery, and the authors extract 40 interpretable rules from search trajectories. The idea of using structured, certifiable feedback to automate scientific model discovery is conceptually aligned with your interest in agentic AI for particle physics tasks, even though the application domain (neutrino flavor models) is outside your primary focus.

Until the next announcement!

P.S. In the most recent submission window: hep-ph: 27 papers hep-ex: 15 papers hep-th: 21 papers