Hey folks,

Here's what you might find interesting today.

Jet Tagging and Jet Physics with ML

Assessing Parameter Redundancy in Transformers for Jet Tagging This might be interesting for you, Aritra. The paper introduces an hourglass architecture that replaces the feed-forward networks inside attention blocks of the Particle Transformer (ParT) and MIParT, while preserving the particle-interaction attention mechanism. The result is roughly 48% and 39.7% parameter reduction for ParT-HG and MIParT-HG respectively, with less than 1% degradation in AUC on the JetClass dataset and comparable background rejection at fixed signal efficiencies for top tagging and quark-gluon discrimination. The key insight is that the FFN blocks carry significant parameter redundancy that can be compressed without losing the physics-aware attention structure. This matters because lighter transformer taggers are more deployable in real-time trigger environments, and the approach isolates which architectural components are actually driving the discriminating power.

Functional anatomy of Pythia-Herwig differences with Kolmogorov-Arnold networks This is interesting for you, Aritra, but you might also find it interesting, Jingjing. The paper uses Kolmogorov-Arnold networks (KANs) to decompose the classifier-derived log density ratio between Pythia and Herwig into explicit one-dimensional observable responses, which can be isolated, recomposed, and transported between generator stages (shower-only, hadronized, full-generator). The authors find that the Pythia-Herwig difference is driven mainly by multiplicity at shower level, shifts toward jet mass and shape after hadronization, and develops a mixed structure in the full generator. Critically, they show that shower-level multiplicity information retains reweighting power downstream, whereas shape responses do not necessarily persist. For Aritra, this is a novel information-theoretic dissection of jet physics through ML; for Jingjing, the explicit decomposition of a learned density ratio into interpretable components is directly relevant to density ratio estimation methodology.

Jet Substructure and Precision QCD

Energy Correlators in V + X as a Benchmark Observable for Precision QCD This might catch your eye, Aritra. The paper proposes projected energy correlators measured on the recoiling QCD radiation of a Z or photon as a benchmark observable, deliberately avoiding jet algorithms by using the boson as the hard scale. The authors present the first complete NLO+NNLL matched predictions for energy correlators at the LHC, combining fixed-order amplitudes, high-order resummation, and universal non-perturbative matrix elements. While not an ML paper per se, this is directly relevant to your interest in jet substructure physics: energy correlators are among the most theoretically clean substructure observables, and this framework provides a precision benchmark that any ML-based substructure method should be compared against. The path to NNLO outlined here could set the theoretical floor for future ML-assisted measurements.

ML for Detector Reconstruction

Statistical validation of calorimeter inpainting with generative diffusion priors This might be interesting for you, Jingjing. The paper tackles localized detector inefficiencies in calorimeter data by using pretrained diffusion models as Bayesian priors to reconstruct missing signal conditioned on surrounding measurements. The authors conduct a systematic comparison of several diffusion-based inpainting algorithms, evaluating them with Bayesian posterior diagnostics for energy response, spatial bias, and uncertainty calibration across collision centralities and masked region sizes. This is relevant to your interest in ML uncertainty quantification: the Bayesian framing provides principled uncertainty estimates on the reconstruction, and the validation strategy (posterior diagnostics rather than point-estimate metrics) is a methodology that could transfer to other generative-model applications in HEP. The work is set in the context of heavy-ion collisions but the approach is general.

Until the next announcement!

P.S. In the most recent submission window: hep-ph: 42 papers hep-ex: 24 papers hep-th: 51 papers