Hey there,

Here's what you folks might find interesting today.

Generative Models and Uncertainty Quantification for LHC Physics

Know What You Don't Flow by Anja Butter, Sascha Diefenbacher, **Tilman Plehn**, and Lorenz Vogel tackles a critical gap in generative ML for LHC physics: calibrated uncertainty estimates. They show how heteroscedastic and Bayesian normalizing flows can learn both systematic and statistical uncertainties on the underlying phase-space density, using a toy model with explicit likelihood first, then extending to a classifier-reweighted approximate generative network when no explicit likelihood is available. The demonstration on top-pair events with a conditional heteroscedastic flow propagating calibrated uncertainties to all phase-space directions is particularly compelling. This might be interesting for you, Jingjing, as it sits squarely at the intersection of **ML uncertainty quantification** and generative modeling, and the classifier-reweighting approach connects to **density ratio estimation** techniques relevant to simulation-based inference workflows.

Differentiable Parametric Simulation and Reconstruction Models in Parnassus by Abdelrahman Elabd, Eilam Gross, Dmitrii Kobylianskii, Runze Li, and **Benjamin Nachman** bridges the gap between hand-crafted parametric detector simulations (like Delphes) and learned generative models by making the parametric prescriptions fully differentiable. The key innovation is that Delphes-style parameters can now be fit to a target detector sample by gradient descent, yielding interpretable, inexpensive models that run in a GPU-enabled Python pipeline. They demonstrate closure by recovering known generating parameters and present a first fit to CMS full simulation. This is interesting for you, Jingjing, as it directly addresses **workflow automation for MC sample generation** and could streamline simulation pipelines, while the differentiable nature opens doors to end-to-end optimization in **simulation-based inference** setups.

Novel ML Architectures for Jet and Event Reconstruction

Down-Type Jet Identification in Fully Hadronic ttbar Events by Jinhui Guo and Chenhao Peng introduces a two-stage ML reconstruction for extracting down-type quark jets from the challenging all-hadronic ttbar channel. The first stage uses a **GNN+Transformer** to assign reconstructed jets into legal six-jet top-pair candidates, and the second stage combines an assignment scorer with an auxiliary **conditional diffusion** head to resolve down/up quark hypotheses inside hadronic W decays. A notable finding is that the diffusion-based assignment score only works when trained with a margin ranking objective, while standard denoising objectives leave it at random baseline, and the two objectives are largely decoupled. This might be interesting for you, Aritra, as it combines **novel ML architectures** (GNN, Transformer, diffusion) for **jet tagging** and exploits **jet substructure** information in a physically motivated way, with direct implications for spin-correlation measurements.

Scouting Super Resolution: from HLT-Level Reconstruction to Offline-like Quality by Maurizio Pierini addresses a practical but important problem: bringing HLT-level scouting data to offline-quality resolution using a trained regression on object-level features and the surrounding particle cloud. A single architecture, applied unchanged to electrons and jets, suppresses residual biases by one to two orders of magnitude and achieves resolution improvements that bring corrected objects within offline calibration uncertainty bands. This is interesting for you, Jingjing, as it demonstrates how ML can effectively **improve reconstruction** quality in a way that could feed into downstream **simulation-based inference** pipelines, and the particle-cloud regression approach is conceptually related to techniques used in density estimation tasks.

Until the next announcement!

P.S. In the most recent submission window: hep-ph: 54 papers hep-ex: 13 papers hep-th: 47 papers