Hey folks,

Here's what you might find interesting today.

Anomaly Detection and Novel ML Architectures for Collider Physics

Contrastive Learning for Interpretable Anomaly Detection at Collider Experiments -- This might be interesting for you, Aritra. The paper introduces ORCA, a two-stage framework that first uses supervised contrastive learning to build an embedding space across diverse physics processes, then runs an autoencoder in that space for anomaly scoring. The key novelty is interpretability: because known processes occupy distinct regions of the contrastive embedding, a maximum-likelihood template fit can attribute anomalous events to known physics processes with quantified uncertainties, even for signals excluded from training. This addresses two persistent problems in LHC anomaly detection -- uninterpretable scores and spurious correlations with energy scale and object multiplicity. The approach is demonstrated on HL-LHC conditions and shows improved breadth and depth of sensitivity over a baseline autoencoder.

Pairton: Iterative Reconstruction of Short-Lived Particles -- This is also relevant for you, Aritra. Pairton formulates particle reconstruction as a masked prediction problem over graph structures, using a pairformer-based architecture with dynamically updated pairwise representations to iteratively predict edges in the decay adjacency matrix. The method learns conditional distributions consistent with a factorised decomposition of decay products, incorporating global event consistency in a way that standard graph neural networks for jet tagging typically do not. It achieves state-of-the-art performance on fully hadronic ttbar decays and is designed to be generalisable to other topologies. The connection between generative modelling concepts (masked prediction) and particle reconstruction is a fresh angle on the graph-based ML paradigm for collider physics.

ML for Inference and Distribution Reconstruction

Neural network maximum entropy framework for distribution reconstruction in heavy-ion collisions -- You might find this tangentially relevant, Jingjing. The paper develops a neural-network maximum-entropy framework for reconstructing probability distributions from limited observables, combining flexible NN representations with Shannon-entropy regularization. The method is applied to extracting conditional jet energy-loss distributions from single-inclusive jet R_AA data, which connects to ML-based unfolding and deconvolution problems. While the application is heavy-ion collisions rather than LHC SBI, the core idea of reconstructing underlying distributions from observables through differentiable forward maps is conceptually adjacent to simulation-based inference and density estimation techniques you work with.

Bayesian inference of event-by-event collision geometry from charged-particle multiplicity in heavy-ion collisions -- This one is a stretch for you, Jingjing, but worth a glance. The paper proposes a Bayesian framework that replaces hard-cut centrality classification with probabilistic assignment based on N_part, using a Monte Carlo Glauber model prior and negative binomial likelihoods. While this is traditional Bayesian inference rather than neural simulation-based inference, the methodology of propagating posterior uncertainties through to downstream observables (net-proton cumulants) and the closure-test validation strategy are relevant to the broader conversation about uncertainty quantification in inference pipelines.

Until the next announcement!

P.S. In the most recent submission window: hep-ph: 24 papers hep-ex: 7 papers hep-th: 23 papers