Hey there,

Here's what you folks might find interesting today.

Jet Tagging and Novel ML Architectures

JetCoRD: Reliability-Aware Cross-Experiment Distillation of Jet Taggers with Adaptive Corrective Representation

This might be interesting for you, Aritra. The paper introduces the first cross-experiment knowledge distillation framework for jet taggers, compressing ATLAS GN2 (5M params) and CMS ParT (2M params) into a single 82k-parameter student. The key innovation is a per-sample reliability signal that simultaneously weights the distillation loss, controls prototype-based teacher repair, and anchors an inference-time gate mixing teacher and student logits. The student actually *exceeds* both teachers at physics-actionable working points -- for instance, +4.3% on b-vs-c at 77% efficiency. The reliability-coupled design is a genuinely new idea in HEP distillation and directly addresses the problem that teacher taggers are imperfect at high-purity working points, which is often glossed over.

Mitigating Detector Ageing Effects with Graph-Based Multi-Modal Track Reconstruction at Belle II

This one is also for you, Aritra. While not about jets per se, it tackles a problem relevant to anyone working with graph neural networks at colliders: how do GNN-based reconstruction algorithms degrade under realistic, time-dependent detector conditions? The authors formulate track finding as a global relational clustering problem using object condensation and show that detector ageing can be treated as a domain shift rather than requiring a fundamentally new strategy. After retraining on degraded conditions, the GNN limits absolute efficiency loss to 14% versus 28% for the Belle II baseline, while maintaining 96% purity. The framing of detector degradation as a domain-shift problem is a useful conceptual lens that could transfer to jet tagging robustness studies.

Generative Models for Collider Physics

Learning Standard Model structure from LHC data with Riemannian flow matching

This is interesting for you, Jingjing. The authors design ShellFlow, a Riemannian conditional flow matching model that generates each particle on its on-shell manifold, trained on ~10^9 real ATLAS Open Data events with no physics input beyond the on-shell condition and invariant-mass formula. From a single training run, the model reproduces dilepton resonances at their PDG positions, the leptonic Weinberg angle, and the W and top masses -- quantities that never enter the training objective. While this is a generative modeling paper rather than simulation-based inference per se, the flow matching methodology and the demonstration that structured manifold constraints can dramatically improve generative quality are directly relevant to your interest in ML techniques for density estimation and simulation in particle physics. The fact that inter-particle correlations are learned without being in the objective is a striking result worth understanding.

That's all that caught my eye today -- a slim but high-quality batch.

Until the next announcement!

P.S. In the most recent submission window: hep-ph: 25 papers hep-ex: 4 papers hep-th: 26 papers