Hi all,

Here's what you folks might find interesting today.

Anomaly detection, jet substructure, and collider ML

[Dissecting Parton Showers with Multi-Point Energy Correlators]
This paper looks especially relevant for you, Aritra. It studies **four-point energy correlators** as experimentally measurable probes of jet substructure, with a clear focus on separating different factorization channels and disentangling intrinsic spin correlations from purely kinematic azimuthal effects. What makes it useful is that it is not just formal substructure theory: the authors compare **Herwig and Pythia** against analytic expectations in controlled limits and also provide parameterizations and analysis algorithms that could be used at the LHC. That is very much in the direction of using refined observables to understand what parton showers are really getting right or wrong inside jets. There is also an **author boost** here: **Philip Harris** is on the paper. Even though this is not an ML paper, it is squarely in your priority area of using sophisticated tools to probe the physics of jets more deeply.

[Transformer-based machine learning using low-level calorimeter signals for collimated photon identification at collider experiments]
This might be interesting for you, Aritra, and there is also a weaker connection for you, Jingjing, because of the trigger-level and analysis-method angle. The main novelty is the direct use of **cell-level calorimeter information** to distinguish highly collimated diphotons from single photons, with a **Transformer** outperforming six compared architectures and an **MLP-Mixer** offered as a lighter alternative for real-time deployment. What stands out is that the paper goes beyond classification and shows **invariant-mass regression directly from calorimeter cells**, which is a concrete physics-facing use of low-level representation learning rather than a generic benchmark exercise. For Aritra, this fits the broader theme of novel ML architectures for collider object identification and of extracting more structure from detector-level information than conventional high-level variables allow. It is not about jets specifically, but it is close enough in spirit to modern low-level collider ML that I would keep it on the short list.

Simulation, reweighting, and analysis infrastructure

[Optimal-Transport-Based Cell Resampling for Negative and Pathological Event Weights]
This is probably the best match today for you, Jingjing. The paper tackles a very practical but important bottleneck in LHC workflows: **negative and pathological MC event weights**, and proposes cell-resampling methods built on **optimal transport metrics** such as the **Energy Mover's Distance** and a spectral variant. The interesting part is that these metrics are soft/collinear insensitive, so the method can be applied directly at different stages of event generation, and the authors report reduced bias relative to earlier resampling approaches on NLO QCD samples. They also introduce the **Cross-Section Mover's Distance** as an unbinned figure of merit for evaluating bias from full-phase-space reweighting, which feels especially relevant if you care about robust density-ratio style comparisons or simulation-based workflows. This is not simulation-based inference in the narrow sense, but it is very much in the orbit of **analysis automation, sample handling, and statistically principled reweighting** for collider studies. I would flag this as worth a careful read because it addresses a pain point that often limits downstream ML analyses more than the inference model itself.

Quantum computing for collider physics

[Overview of Applications of Quantum Computing in QCD]
This one is more peripheral, but I think it is still worth flagging for you, Aritra. It is a concise overview of **quantum computing applications in collider/QCD problems**, covering event generation, parton showers, amplitudes, loop and phase-space integration, and optimization tasks relevant to experiment. The reason it makes the list at all is your interest in **Lorentz-equivariant architectures including for quantum circuits**, although this paper is a broad review and does **not** appear to discuss equivariant quantum ML architectures specifically. So I would not treat it as a must-read technical paper, but rather as a quick landscape update if you want to keep tabs on

where quantum methods are being proposed in HEP. For Jingjing, the relevance is weaker, since there is no clear SBI, uncertainty, or workflow-automation component in the abstract.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 32 papers hep-ex: 12 papers hep-th: 39 papers