Hi all,

Here's what you folks might find interesting today.

TOP PICKS

Anomaly detection, triggers, and agentic AI for collider workflows

- **Learning to Trigger: Reinforcement Learning at the Large Hadron Collider**
  This paper turns online trigger-threshold tuning into a sequential decision problem and uses reinforcement learning to adapt thresholds as detector conditions and pileup drift over time. What makes it stand out is that they do not just test on simulation: they report transfer to *real CMS collision data* without fine-tuning, including an anomaly-detection trigger based on reconstruction loss. For you, Aritra, that AD-trigger angle is the key hook, since it connects directly to model-agnostic rare-signal searches under realistic operational constraints rather than just offline anomaly scoring. For you, Jingjing, this is also relevant as a concrete example of *agentic AI for analysis operations*, where the AI is not classifying events once, but actively controlling part of the experimental workflow. It is also worth noting that **Jennifer Ngadiuba** is on the author list, which is a strong signal that this is plugged into serious LHC-ML practice.

- **One Generator, Any Process: LLM-Conditioning for the LHC**
  The central idea here is to train a single generative model across many LHC processes by conditioning on continuous parameters, process labels, and even Feynman diagrams, with embeddings supplied by pre-trained LLMs. That is more ambitious than standard per-process surrogate generation: they are trying to inject high-level physics structure so the model can generalize to *unseen* processes rather than just interpolate within one channel. This might be interesting for you, Jingjing, because it sits right at the intersection of workflow automation and simulation infrastructure, and could matter for faster sample generation or more reusable ML-based simulators. Aritra, you may also want to keep an eye on it because foundation-model conditioning for event generation could eventually feed into anomaly-search pipelines by broadening the space of efficiently modeled backgrounds and signals. Also, **Tilman Plehn** is on the paper, so this gets an author boost on top of the topical relevance.

Jet tagging and jet substructure

- **Application of Deep Learning to Jet Charge Discrimination**
  This is the most directly aligned jet-ML paper in today's list: they benchmark several classical and quantum ML approaches for distinguishing up-quark from anti-up-quark jets, with a GNN coming out on top. For you, Aritra, the relevance is straightforward because this is a jet-tagging problem that goes beyond generic quark/gluon separation and focuses on charge-sensitive information in the radiation pattern. The useful part is not just the headline AUC, but the comparative setup across model classes, including quantum ML, which could help calibrate how much architectural sophistication is actually buying you on a controlled jet task. It also has clear phenomenological motivation, since quark-versus-antiquark discrimination can feed into asymmetry measurements and BSM searches. If you are thinking about where novel architectures genuinely exploit jet substructure rather than just repackage standard features, this is probably the jet paper to scan first today.

- **Mapping jet substructure in heavy-ion collisions with track functions**
  This one is not an ML paper, but it is still likely worth your attention, Aritra, because it studies jet substructure through *track functions*, which are sensitive to the full fragmentation pattern and obey nonlinear RG evolution. The novelty is that they use higher moments and cumulants of these track-function distributions to expose medium-induced modifications and to discriminate between JEWEL and HYBRID quenching pictures. That makes it more than another heavy-ion phenomenology note: it is proposing a structured observable family that captures detailed information flow inside jets. If you are interested in ML for understanding jet physics, this kind of observable design can be exactly the sort

of physics prior that later becomes useful input, target, or inductive bias for learning-based analyses. I would treat it as a good "physics of jets" paper rather than a methods paper.

- **Event isotropy in perturbative QCD**
  This paper develops the theory of *event isotropy*, an observable built from the Energy Mover's Distance that quantifies how close an event is to a uniform energy distribution. That should be interesting for you, Aritra, because EMD-based geometric descriptions of collider events have repeatedly shown up in ML-adjacent work on anomaly detection, representation learning, and event-level structure. The novelty here is that they push the observable into a proper perturbative-QCD setting, with fixed-order results and NLL resummation, so it becomes something theoretically controlled rather than just geometrically appealing. In practice, that matters if one wants to use isotropy-like quantities in searches or weakly supervised pipelines without losing contact with calculability. It is not a direct ML paper, but it looks like a useful bridge between modern event geometry and robust collider observables.

Simulation-based inference, information geometry, and broadly useful ML/statistics

- **The Degeneracy Distillery**
  This is probably the strongest general-methods paper for you, Jingjing. They propose a way to automatically detect and symbolically resolve parameter degeneracies from parameter-simulation pairs by estimating and flattening the Fisher information, then use the resulting coordinates to reduce the simulation cost of downstream neural posterior estimation. That is highly relevant to simulation-based inference because degeneracies are one of the main reasons SBI pipelines become sample-hungry, unstable, or hard to interpret in high-dimensional physics problems. The especially attractive claim is that the transformed coordinates flatten the Fisher information *globally in expectation*, rather than only locally around one posterior mode. Even though the paper is not explicitly particle-physics focused, it is exactly the kind of information-geometric tool that could transfer well to likelihood-free inference workflows.

- **Data Augmentation: A Fourier Analysis Perspective**
  This is a theory paper on when *partial* data augmentation can recover the same statistical gains as full group augmentation, using Fourier analysis and representation theory of finite groups. Jingjing, this may be useful for you as a general ML methods read because many particle-physics pipelines rely on approximate symmetry enforcement, and this paper gives a principled way to think about the tradeoff between exact invariance and computational cost. Aritra, you might also find it relevant if you are thinking about symmetry-aware architectures or augmentation strategies for jets and event representations. The nice part is that it does not just say "augmentation helps"; it characterizes when a random subset of group actions is enough and when exact invariance is impossible without averaging over the full group. That kind of result can inform practical design choices in physics ML where symmetry is important but full equivariance is expensive.

Quantum and many-body ML with possible HEP crossover

- **Fully-heavy multiquarks in neural-network quantum states**
  This paper applies neural-network quantum states to fully-heavy multiquark spectroscopy, representing the many-body spatial wavefunction with deep nets while enforcing the color-spin structure exactly from group theory. It is not directly targeted at either of your hottest topics, but it is one of the more plausible crossovers from modern ML into hadron physics in today's batch. The interesting technical point is that they use NNQS to tackle the high-dimensional correlation structure that makes traditional few-body methods expensive or inflexible. For you, Aritra, this may be worth a skim if your interest in quantum/ML methods includes architectures that respect physics constraints rather than generic black-box fitting. I would rank it below the trigger, generator, and jet papers, but it is a credible "methods migrating into HEP theory" pick.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 38 papers hep-ex: 8 papers stat.ML: 11 papers