Hi all,

Here's what you folks might find interesting today.

TOP PICKS

Jet physics and substructure

- **Putting Jet Substructure on Track(s)**
  This is probably the clearest jet-focused pick today for your interests. The paper is not an ML paper, but it pushes the theory side of **jet substructure** in a way that could matter a lot for downstream ML work, especially for architectures that try to exploit fine-grained angular information. The main novelty is the first complete calculation of track-based jet substructure observables at the LHC, including projected energy correlators up to four points at next-to-leading collinear logarithmic accuracy. That matters because tracking gives much better angular resolution than calorimetry, so this sharpens the physics targets that jet-tagging or representation-learning methods should be trying to capture. This might be interesting for you, Aritra, since it is directly about understanding jets more deeply rather than just improving classifier benchmarks. Even without an ML component, it looks like useful reading for grounding future work on substructure-aware ML.

- **Gluon radiation from a QCD antenna with realistic parton-medium interactions**
  This one is also more on the theory side, but it is relevant to the broader theme of **understanding the physics of jets**. The authors replace common approximations like the harmonic-oscillator limit with a numerical differential-equation approach that resums medium interactions to all orders using more realistic scattering models. The payoff is a more faithful picture of angle and energy distributions for emitted gluons, and of how color coherence breaks down across phase space. This might be interesting for you, Aritra, because it gives a more detailed physical picture of radiation patterns that could eventually inform ML observables or inductive biases for jet studies in heavy-ion or medium-modified settings. It is not directly about tagging, but it does look useful if you are thinking about what structure a model should learn from first principles.

LHC ML for inference and event characterization

- **Higher-order effects in amplitude-assisted polarisation extraction with machine-learning techniques**
  This is one of the strongest ML-relevant papers in today's list. The key novelty is an **amplitude-assisted regression** setup pushed to next-to-leading-order QCD accuracy and matched to parton showers, rather than staying at a simplified truth-level or leading-order benchmark. They compare several neural-network architectures against a random-forest regressor for extracting the longitudinal gauge-boson production rate in diboson events, so the paper is really about how ML behaves once realistic theory systematics are included. This might be interesting for you, Jingjing, because it sits near the boundary between precision inference and ML-based extraction of latent physics quantities from collider data. Aritra, you might also find it useful as an example of ML integrated with realistic LHC event modeling rather than just black-box classification. While it is not simulation-based inference in the strict neural-posterior sense, it is a serious example of ML-assisted parameter-sensitive event interpretation.

- **Modeling Falling Backgrounds with Exponential Mixtures**
  This is a statistics-heavy collider methods paper rather than an ML paper, but it feels quite relevant to modern search workflows. The authors propose **finite exponential mixtures** as a semi-parametric model for smoothly falling backgrounds, motivated by extreme-value theory, and test them on both very large and small published datasets. What stands out is that they explicitly study bias and coverage, finding small bias relative to statistical uncertainty while maintaining nominal coverage in simulation. This might be interesting for you, Jingjing, because it connects to robust inference and uncertainty-aware modeling in

search analyses, especially where background flexibility and calibration matter more than raw classifier performance. Aritra, there is also some relevance here if you are thinking about model-agnostic search strategies, since anomaly-style searches often live or die by how backgrounds are parameterized. It is not flashy, but it looks practically useful.

Quantum information and collider-adjacent quantum structure

- **Quantum entanglement of photon pairs at proton-proton colliders**
  This is a neat collider-quantum-information crossover paper. The novel idea is to use the Bethe-Heitler conversion process in the tracker as an effective spin analyzer for photons, so that the angular distributions of the resulting electron-positron pairs encode the diphoton polarization state and hence entanglement information. The authors then estimate HL-LHC sensitivity, finding that Higgs to gamma gamma is far too statistics-limited for a convincing entanglement measurement, while continuum diphotons could reach about 1.5 sigma. This might be interesting for you, Aritra, especially because you have an interest in Lorentz-aware and quantum-structure questions around collider systems, even though this is not a quantum-ML paper. The significance projections are modest, but the measurement concept itself is the real reason to read it.

- **Production of Magic States via Z Bosons and Dark Photons**
  This is more speculative and less directly tied to your core priorities, but it does sit in the quantum-information-meets-HEP space. The paper studies how scattering in the electroweak sector and in a dark-photon extension can generate **magic states**, tracking how the distribution over stabilizer classes changes across energy regimes and mediator structures. For you, Aritra, the relevance is mainly conceptual: it is another example of importing quantum-information diagnostics into particle processes, though not yet in an ML or collider-analysis framework. I would rank it below the diphoton entanglement paper because the path to phenomenological use is less clear. Still, if you are scanning for quantum-information ideas that might eventually connect to collider observables or quantum architectures, it may be worth a look.

Broader ML for fundamental theory and workflow search

- **Exploring Line Bundle Standard Models with Transformers**
  This is outside collider phenomenology, but it is one of the more substantive ML-for-fundamental-physics papers in the list. The authors build a Transformer-based reinforcement-learning system, "LB-Explorer", to search heterotic line-bundle Standard Model constructions on Calabi-Yau compactifications, and then augment it with CP-SAT solvers to impose some constraints exactly. The interesting part is not just "using transformers", but the hybrid search strategy: learned exploration for a huge constrained landscape, followed by exact combinatorial refinement. This might be interesting for you, Jingjing, because it touches the broader theme of **agentic or workflow-oriented AI for physics tasks**, even though the application is string model building rather than collider analysis. I would not call it a top-priority read for your immediate projects, but it is a good example of ML being used as a structured scientific search tool rather than a classifier.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 34 papers hep-ex: 12 papers hep-th: 25 papers