Hi all,

Here's what you folks might find interesting today.

TOP PICKS

Jet ML and event representations

[Machine learning fully hadronic events with spectral functions]

This one looks like the clearest match today for **Aritra**. The key idea is to represent a fully hadronic event by a **two-point correlation spectral function**, a 1D object in angular distance that is invariant to collider isometries and jet permutations and, importantly, does not depend on a fixed jet multiplicity. That directly targets a real pain point in hadronic LHC ML, where ISR/FSR changes the number of jets and makes fixed-input architectures awkward. They then feed these spectral-function features into a dense network for a gluino-to-ttbar+MET benchmark and report a sizeable gain in expected reach relative both to a recent ATLAS analysis and to the same network trained only on jet kinematics. This is not anomaly detection, but it is very much in the lane of **ML for jets / hadronic structure**, with a representation that compresses combinatorial event information in a physics-aware way. If you are thinking about alternative event encodings for busy hadronic final states, this seems worth a close read, Aritra.

Quantum ML for detector or collider data

Qudit extension of parameterized IQP circuits: A generative quantum machine learning approach to integer data

This feels like a more speculative but still relevant pick for **Aritra**, especially given the interest in **quantum ML architectures**. The paper extends parameterized IQP circuits from qubits to **qudits** so that integer-valued detector data can be modeled without first forcing it into binary encodings that distort the original metric structure. Their test case is calorimeter shower data from the CLIC electromagnetic calorimeter, and they also develop the associated training loss and feature covariance machinery for this non-binary generative setting. What makes it potentially useful is not just "quantum ML on HEP data", but the specific architectural move: preserving the geometry of integer-valued observables rather than flattening them into bit strings. That could matter for any future attempt to build quantum generative models for richer detector-level or jet-level observables. It is not Lorentz-equivariant, so it is not a perfect match, but it is one of the better quantum-ML-adjacent HEP papers in today's list.

Simulation-based and differentiable inference methods

HIcosmo: a differentiable JAX-based framework for cosmology inference

This is the best general-methods paper for **Jingjing** today, even though it is in cosmology rather than collider physics. The authors build a fully **differentiable JAX-based inference stack** in which the forward model, likelihood, posterior sampling, and Fisher forecasts are all written in JAX primitives, so gradients and Hessians come directly from autodiff rather than finite differences. They validate against Cobaya and report substantial throughput gains, including strong GPU acceleration as data volume grows. The reason this may be worth your attention, Jingjing, is that the paper is really about **end-to-end differentiable statistical inference infrastructure**, which is closely aligned with the broader SBI / likelihood-free / high-dimensional inference ecosystem even if the implementation here is likelihood-based. If you are thinking about scalable inference pipelines for particle physics, the design choices around autodiff-native modeling and hardware acceleration could transfer well. Not a direct SBI paper, but definitely adjacent in a useful way.

NN parameterizations for hadron structure

Neural-Network extraction of TMDs with SIDIS data

This is another methods-adjacent paper that may be of interest to **Aritra**, and possibly also **Jingjing** from the inference angle. The author performs a first global extraction of unpolarized **TMDs** using a **neural-network parameterization**, combining DY and SIDIS data at very high perturbative accuracy. The notable point is that the NN parameterization is being used to reduce model dependence in the functional form, and the inclusion of SIDIS broadens the extracted TMDs while shrinking uncertainties relative to a DY-only fit. This is not collider ML in the LHC-analysis sense, but it is a concrete example of ML being used to learn nonperturbative QCD structure from data rather than imposing a rigid ansatz. For you, Aritra, the appeal is the connection to **jet/QCD structure and information extraction from hadronic observables**; for you, Jingjing, it is a reminder of how flexible learned parameterizations can change inference systematics. More niche than the spectral-function paper, but still plausibly worth skimming.

Information-theoretic observables and QCD structure

The one-point charge correlator in deep inelastic scattering

This is not an ML paper, but it is probably the strongest **information-theoretic / substructure-style observable** paper in the list, so I think it is worth flagging for **Aritra**. The authors define a new **one-point charge correlator** in DIS, show that it is IRC safe, and analyze it in SCET in both forward and back-to-back limits. What is especially notable is that in one limit it introduces a new nonperturbative object encoding multidimensional nucleon structure, while in the back-to-back limit it connects directly to standard TMD factorization. That kind of observable-building is often where future ML applications become interesting: you get a theoretically controlled summary statistic with clear physics content, rather than a black-box feature space. So while this is not immediately usable as an ML architecture paper, it looks relevant to the broader theme of **learning or exploiting fine-grained hadronic structure**. If you are thinking about physics-aware features for QCD-rich problems, this seems like a good conceptual paper to keep on the radar.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 41 papers hep-ex: 9 papers hep-th: 43 papers