Hi all,

Here's what you folks might find interesting today.

TOP PICKS

Jet physics and ML for substructure

Disentangling Dark Gauge Symmetries with Deep Learning on the Lund Jet Plane

This paper is the clearest jet-ML match in today's list. The authors build a new Monte Carlo shower framework for arbitrary dark gauge groups, with explicit handling of generalized color topologies, recoil, and full mass effects through an exact three-body phase-space treatment. On top of that, they analyze the resulting radiation patterns in the **Lund Jet Plane** using a **Neural Sorter Mamba Network**, so the ML is tied directly to structured jet-emission information rather than generic event-level features. This might be interesting for you, **Aritra**, because it is really about learning physically meaningful differences in radiation structure, and the Lund-plane angle makes it close in spirit to jet substructure representation learning. Even though the target is dark-sector gauge discrimination rather than standard jet tagging, the combination of new shower modeling plus sequence-style ML on substructure patterns feels transferable to LHC jet studies. The novelty here is not just "deep learning on jets", but a pipeline where the simulation itself is designed to expose symmetry-dependent emission topology that the network can exploit.

In-medium QCD splittings beyond the soft, large-N_c and harmonic-oscillator approximations all at once

This is not an ML paper, but it is still worth flagging because it pushes the theory side of jet structure in a way that could matter for ML-driven jet studies. The authors report what they describe as the first complete numerical solution of the **BDMPS-Z** equations, simultaneously going beyond soft, large-N_c, and harmonic-oscillator approximations. That means more realistic in-medium splitting functions across phase space, including finite-energy and subleading-color effects, which is exactly the kind of improved physics input that can sharpen downstream jet-observable design or ML-based heavy-ion analyses. This might be interesting for you, **Aritra**, because it is fundamentally about understanding the physics of radiation patterns and where standard approximations fail. It is not directly aligned with anomaly detection or tagging architectures, so I would treat it as a lower-priority but still relevant theory-side read.

Simulation, inference infrastructure, and analysis automation

Parnassus: A GPU-enabled, Python-based Package for Fast Particle Detector Simulation and Reconstruction

This is probably the broadest practical methods paper in the list. The package provides a **Python/PyTorch**, GPU-capable framework for fast detector simulation and reconstruction, with interchangeable backends ranging from parametric Delphes-style responses to **flow-matching neural detector models**. The process-agnostic detector-card interface is a particularly useful design choice: the same released detector model can be applied to new processes without retraining, which makes it much more analysis-friendly than many bespoke ML simulators. This is interesting for you, **Jingjing**, because it sits right at the interface of modern ML methods and workflow acceleration for collider analyses, and it could be useful infrastructure for simulation-based studies or automated analysis pipelines. This is also worth a look for you, **Aritra**, especially if you are thinking about fast turnarounds for jet-based studies or anomaly-search prototyping. Also, **Benjamin Nachman** is on this paper, which is an extra reason to pay attention given your interests.

Di-Higgs and Higgs phenomenology

Electroweak corrections to Higgs boson pair production: The quark channel

This is the main **di-Higgs** paper in today's set, though it is a phenomenology calculation rather than an ML-driven analysis. The authors compute mixed QCD-electroweak corrections to Higgs-pair production in the quark-antiquark channel, with fully analytic virtual amplitudes obtained via differential equations and then implemented in **POWHEG-BOX**. The notable result is that these corrections can distort differential distributions by up to about 10 percent near threshold in the HH invariant mass spectrum, which is exactly the kind of effect that can matter for precision projections and shape-based searches. This might be interesting for you, **Jingjing**, because even without an explicit ML component, it updates the theory systematics for HH studies and could affect how one thinks about training targets or signal modeling in future analyses. I would rank it below an HH+ML paper, but among today's options it is still the most directly relevant di-Higgs item.

General ML and statistics methods with possible HEP relevance

Statistically Valid Hyperparameter Selection: From Tuning to Guarantees

This is a methods paper rather than a particle-physics one, but it is unusually relevant to anyone deploying ML in high-stakes analyses. The monograph frames hyperparameter selection through the **learn-then-test (LTT)** paradigm, treating tuning as a multiple-testing problem and giving finite-sample guarantees for reliability constraints such as average risk, quantile risk, and information-theoretic criteria. That is potentially useful whenever one wants principled control over threshold choices, calibration settings, or inference-time knobs instead of relying on ad hoc validation sweeps. This might be interesting for you, **Jingjing**, especially for simulation-based inference pipelines or automated workflows where one would like guarantees on selected settings. It may also be worth a skim for you, **Aritra**, if you are thinking about robust operating-point selection in anomaly detection or weakly supervised searches.

Information from coincidences

This is a fairly abstract information-theory paper, but it unifies a surprisingly broad set of variational identities under one algebraic "mixed coincidence" formula. The paper claims a common derivation of results touching Sanov-type decompositions, Chernoff information, Donsker-Varadhan, PAC-Bayes, and several rare-event coincidence laws, with extensions to multi-prior settings. For your purposes, the most relevant angle is that it may provide a compact conceptual toolkit for thinking about information-theoretic objectives, change-of-measure arguments, and multi-prior penalties that sometimes show up in ML-for-physics formulations. This might be interesting for you, **Aritra**, because you have an explicit interest in information-theoretic tools for jet physics, even though the paper itself is not about HEP. I would treat this as a speculative theory-methods read rather than an immediate application paper.

Efficient Adaptive Data Acquisition via Pretrained Belief Representations

This one proposes **POLAR**, a framework for adaptive data acquisition that separates belief-state representation learning from policy learning by using pretrained predictive models as belief encoders. The same setup is cast as a unified amortized policy-learning approach for Bayesian experimental design, Bayesian optimization, and active learning, with the task-specific utility defining the downstream objective. The practical claim is improved sample efficiency and scalability relative to existing amortized acquisition methods. This might be interesting for you, **Jingjing**, because the belief-state viewpoint is close in spirit to sequential experimental design and could be relevant to simulation-based or agentic analysis workflows where one adaptively chooses what to simulate or evaluate next. It is not particle-physics specific, but the abstraction level is close enough to modern analysis automation that I think it is worth knowing about.

Gaussian Mean Field Variational Inference can Overestimate Predictive Variance

This paper makes a useful corrective point about a very common approximation: while **MFVI** is usually said to underestimate posterior variance, the authors show that it can actually **overestimate predictive variance** in directions where the training data concentrate. They derive this in conjugate Bayesian linear regression, connect it to the **cold posterior effect**, and argue that temperature scaling can partially fix the mismatch. This might be interesting for you, **Jingjing**, because uncertainty quantification is central in likelihood-free and simulation-based workflows, and this result is a reminder that parameter-space intuition does not automatically transfer to predictive-space behavior. The value here is not a new HEP application, but a sharper understanding of when a standard approximate Bayesian tool can mislead downstream decisions. If you are comparing posterior surrogates or calibrating predictive uncertainties, this is the kind of caveat paper worth having in mind.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 33 papers hep-ex: 14 papers stat.ML: 13 papers