Hi all,

Here's what you folks might find interesting today.

TOP PICKS

General ML methods with likely particle-physics crossover

Interpretability and explainability in physics ML

- **Interpreting "Interpretability" and Explaining "Explainability" in Machine Learning in Physics**
  This is a concepts-and-practice paper that tries to cleanly separate **interpretability** from **explainability**, which is actually useful if you are building ML systems for physics rather than just benchmarking them. The authors define interpretability in terms of structural transparency of the model, and explainability in terms of mapping learned behavior back onto scientific knowledge, then discuss the tradeoffs each imposes on model design. That framing could be especially useful when deciding how much architectural bias to build into anomaly detection or jet-tagging pipelines, versus how much flexibility to leave for discovery. This might be interesting for you, **Aritra**, because it gives a sharper language for thinking about whether a jet model is learning robust substructure physics or just a performant proxy. It may also be worth a look for you, **Jingjing**, since the emphasis on task specification and intervention plans connects naturally to simulation-based inference workflows where model assumptions need to be explicit. **Jesse Thaler** is on this paper, so it gets an author boost as well.

Symmetry-aware ML for geometry and physics

- **The Sharp Edges of Calabi-Yau Manifolds: Designing Symmetric Models for Ricci-flat Metrics**
  This is not particle-phenomenology-facing, but it does hit the "novel ML methods with possible HEP application" criterion quite well. The main technical angle is a **symmetry-aware graph neural network** for approximating Ricci-flat metrics, together with a discussion of how broken symmetries in sampling can spoil learning. That combination of architectural symmetry handling plus data-generation pathology analysis is the part that feels portable beyond string geometry. This might be interesting for you, **Aritra**, especially because you care about architectures that encode physics structure; while this is not Lorentz equivariance, it is very much in the same spirit of building the right inductive bias into the model. More broadly, the paper looks useful as a case study in when respecting exact symmetries improves both stability and physical faithfulness.

- **Generating Special Triangulations with Transformers**
  This paper uses **transformers to generate fine, regular, star triangulations** of high-dimensional polytopes, which is a fairly nontrivial structured-generation problem rather than a standard regression or classification task. The notable point is that the model is not just fitting labels but learning a constrained combinatorial object, and the authors also report a kind of self-improvement loop by retraining on model-generated outputs. That makes it potentially relevant as a methodological example for structured generation problems in physics workflows. This might be interesting for you, **Jingjing**, because agentic or automated analysis pipelines in HEP will likely need this sort of constrained generation and iterative refinement rather than one-shot prediction. It is less directly tied to collider ML, but the workflow idea is probably the transferable part.

Information-theoretic and score-based statistical structure

- **Probing Probability Geometry with Schwinger-Dyson Identities: Score Mismatch, Fisher Information, and Configurational Temperature**
  This is a theory-heavy paper, but the core object is a **score-mismatch field** that controls

departures from an equilibrium target distribution, with Schwinger-Dyson violations interpreted as projections of that field. The reason it may be worth triaging is that it ties together **score functions, Fisher information, Stein operators, and variational characterizations** in one framework. That is the kind of statistical language that can sometimes migrate into likelihood-free inference, goodness-of-fit testing, or diagnostics for learned samplers. This might be interesting for you, **Jingjing**, because simulation-based inference increasingly leans on score-based and information-geometric ideas, even when the application papers use different terminology. It is not a direct SBI paper, but it looks like one of the more plausible theory-to-method crossover items in today's list.

Agentic or RL-assisted exploration of scientific parameter spaces

- **Supercool with PPO: Exploring Supercooled Phase Transitions via Reinforcement Learning**
  This paper applies **proximal policy optimization** to search efficiently through model parameter space for benchmark points with detectable gravitational-wave signals from supercooled phase transitions. The novelty is not just "using RL", but building reward functions around phenomenological objectives like large amplitudes, broad frequency coverage, and detector sensitivity, then comparing directly against conventional Monte Carlo scans. That makes it a concrete example of **goal-directed exploration in high-dimensional theory space**, which is closer to workflow automation than to standard surrogate modeling. This might be interesting for you, **Jingjing**, because it is the kind of paper that hints at how agentic systems could help with benchmark discovery, sample targeting, or adaptive scan design in collider phenomenology too. It is outside your core LHC topics, but the search strategy itself seems transferable.

LHC searches, triggers, and LLP infrastructure

LLP triggering and detector concepts

- **High Multiplicity Trigger for Long-Lived Particles in CMS detector**
  This is an experimental trigger paper rather than an ML paper, but it is directly relevant to **model-agnostic or unusual-signature searches** because it targets displaced LLP decays in the muon system using anomalously high CSC hit multiplicities. The useful part is that it documents the actual deployed trigger logic, pileup dependence, threshold optimization, and Run 3 performance on both simulation and data. For anomaly-search people, papers like this matter because they define what unconventional signatures are realistically collectable online, which in turn constrains downstream ML search strategies. This might be interesting for you, **Aritra**, since broad new-physics searches live or die on triggerability, and this trigger is explicitly designed for non-standard topologies. Even without an ML component, it is a good reality check on the experimental side of LLP-style discovery.

- **Projected sensitivity of the ANUBIS detector to heavy neutral leptons**
  This is a sensitivity study for **ANUBIS** targeting HNLs, with the main practical contribution being a quantified reach and a flexible evaluation framework, **SET-ANUBIS**, for LLP models. It is not especially ML-centric, but it is relevant if you care about the broader landscape of non-standard LHC signatures and where future search opportunities may open up. The paper also notes room for improvement through future analysis strategies, which is the part most likely to intersect with ML later. This might be interesting for you, **Aritra**, as background reading on LLP search infrastructure and complementary detector coverage. I would treat it as lower-priority than the CMS trigger paper, but still potentially useful context.

Other papers I would probably skip today

I would deprioritize the Feynman-integral and KLN-subtraction papers for this particular digest: they look technically strong, but they do not line up closely with your ML-for-particle-physics priorities. Likewise, the Belle II lepton-ID and CMS luminosity papers are important experimentally, but the abstracts do not indicate the kind of ML novelty or direct overlap with your current interests that would push them into top-pick territory.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 33 papers hep-ex: 13 papers hep-th: 40 papers