Hi all,

Here's what you folks might find interesting today.

TOP PICKS

Agentic AI and workflow automation to particle physics analysis tasks

LeWRON: Agentic Analysis of Electroweak Phase Transitions

This paper is probably the clearest match for the workflow-automation thread today. The authors build an *agentic* framework that takes an input Lagrangian and orchestrates the full electroweak phase transition pipeline, including effective-potential construction, thermal history tracking, bubble nucleation, and gravitational-wave predictions. What makes it more than just a wrapper is the audited, artifact-based design: intermediate analytic outputs are checked by auditor agents and stored in a structured way for reproducibility and human inspection. There is also a "reproduction mode" aimed at inferring conventions from the literature and reproducing published results, which is exactly the kind of practical friction point that often slows real phenomenology work. This might be interesting for you, Jing, because it is a concrete example of agentic AI being used for a real particle-physics analysis workflow rather than just code assistance. Even though it is not collider ML in the narrow sense, the automation philosophy feels very aligned with analysis-pipeline acceleration.

Jet tagging with novel ML architectures

Towards Engineering Scaling Laws with Pretraining Data Composition

This one is especially relevant on the jet-ML side because it is not just another jet classifier paper; it asks how the *composition* of pretraining data changes the scaling behavior of large models for hadronic jet classification. The key claim is that, in particle physics, where synthetic data are cheap, one can intentionally steer the scaling regime toward "more data, not necessarily more parameters" by using pretraining samples that are more diverse and better aligned with the downstream task. That is a useful idea if you care about how to build jet-tagging systems efficiently rather than only squeezing out one benchmark number. This might be interesting for you, Aritra, because it directly touches jet classification and model-design strategy, and it may also be useful for you, Jing, given the broader implications for large-model training in HEP. Benjamin Nachman is on this paper, so I would give it an extra look on that basis too. It feels like one of the more conceptually useful papers of the day for anyone thinking about foundation-style models for collider data.

Novel machine learning or statistics methods with possible particle-physics application

HEPTv2: End-to-End Efficient Point Transformer for Charged Particle Reconstruction

This is one of the strongest ML-for-HEP papers in the list, even if it is about tracking rather than jets or inference. The novelty is an end-to-end point-transformer pipeline that reconstructs tracks directly from detector hits without the usual graph-building, clustering, or filtering stages that break end-to-end optimization. The encoder uses locality-sensitive hashing in detector coordinate space to preserve geometry while keeping attention efficient, and the decoder predicts complete trajectories with joint supervision. The reported numbers are impressive not just in accuracy but in the latency-memory tradeoff, which matters a lot for HL-LHC deployment scenarios. This might be interesting for you, Aritra, because it is a serious new architecture for collider reconstruction, and also for you, Jing, as an example of efficient end-to-end ML in a realistic HEP setting. Javier Duarte is on this paper, which is another good signal that it is worth skimming.

Statistical Properties of Training & Generalization

This is a broader methods paper rather than a direct HEP application, but it is unusually well aligned with the kinds of questions that come up when adapting ML to physics. The authors review neural scaling laws and connect them to inductive bias, constraints, and modeling choices in physics-informed settings, which makes it more useful than a generic deep-learning overview. If you are thinking about why certain architectures or training regimes work in collider problems, this kind of perspective can help frame those choices more systematically. This might be interesting for you, Aritra, in the context of architecture choices for jet studies, and for you, Jing, because scaling-law thinking is increasingly relevant for simulation-based and large-model workflows. It is not a paper with a new benchmark result, but it looks like a good "calibration of intuition" read. I would treat it as a background paper with potentially high downstream value.

Calibration without labels in multiple testing

This is a statistics paper, but the core idea could matter for anomaly detection, discovery claims, and any setting where one wants interpretable error probabilities without direct labels. The authors study calibration for large-scale hypothesis testing when ground truth is never observed, and introduce pseudo-labels derived from ordered p-value spacings so that calibration tools can still be brought to bear. That is a pretty elegant conceptual move, especially because many HEP problems live in exactly this awkward regime of partial or absent truth labels. This might be interesting for you, Aritra, if you are thinking about model-agnostic searches or multiple-testing issues in broad scans, and possibly for you, Jing, from the inference-and-uncertainty side. It is not particle-physics-specific, so I would not put it above the direct HEP ML papers, but it has the kind of statistical idea that could travel well. Definitely one for the "methods to keep in mind" pile.

A Solver-Free Training Method for Predict-then-Optimize

This one proposes a decision-focused learning method that avoids calling an optimization solver during training, which is the main bottleneck in many predict-then-optimize pipelines. The technical hook is a measure-transformation-based surrogate loss that is solver-free while still coming with Fisher consistency and excess-risk guarantees. That could be useful anywhere one wants to train ML models against downstream optimization objectives but cannot afford expensive inner loops. This might be interesting for you, Jing, especially if you are thinking about automated analysis pipelines or resource-aware workflow design, and perhaps also for you, Aritra, as a general-purpose ML method worth knowing about. It is not directly particle physics, but it is the sort of optimization-aware training idea that could become relevant in detector operations, scheduling, or analysis design. I would classify it as speculative but potentially high-upside.

ML exploiting jet substructure or understanding the physics of jets

Vistas: A Visualization Interface for Particle Collision Simulations

This is not an ML paper, but I think it still deserves a mention because it offers an interactive way to inspect the full Pythia event-development chain, from hard process through showering, hadronization, and decays. The useful part is that it exposes structures like color flow, beam remnants, and multiple parton interactions in a browser-based 3D interface with stage-by-stage toggling and kinematic filtering. For anyone thinking about jet substructure or trying to build intuition for what ML models might be learning from jets, tools like this can be surprisingly valuable. This might be interesting for you, Aritra, as a qualitative aid for jet-physics understanding, even though it is not itself a learning method. I would not prioritize it over the scaling-laws or HEPTv2 papers, but it looks like a nice practical resource. It could also be handy for teaching or group discussions around event structure.

BY CATEGORY

Anomaly detection at the LHC / Jet tagging / Jet physics / Lorentz-equivariant and related collider ML

- HEPTv2: End-to-End Efficient Point Transformer for Charged Particle Reconstruction
- Towards Engineering Scaling Laws with Pretraining Data Composition
- Vistas: A Visualization Interface for Particle Collision Simulations

Simulation-based inference / Di-Higgs / Agentic AI and workflow automation

- LeWRON: Agentic Analysis of Electroweak Phase Transitions

General ML / statistics methods with possible HEP relevance

- Statistical Properties of Training & Generalization
- Calibration without labels in multiple testing
- A Solver-Free Training Method for Predict-then-Optimize
- Rigorous uncertainty quantification of probabilistic AI weather forecasts with conformal prediction
- AK-MCS-C2 : Active Kriging Monte Carlo Simulation method with conformal certification for failure probability estimation
- Optimal Deterministic Multicalibration and Omniprediction

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 43 papers hep-ex: 18 papers stat.ML: 18 papers