Hi all,

Here's what you folks might find interesting today.

Jet tagging and detector ML

- **Measurement of the b-jet identification efficiency in dileptonic $t\bar{t}$ events using proton-proton collision data at $\sqrt{s}=13.6$ TeV collected with the ATLAS detector**
  This is probably the clearest must-read today for jet-ML interests because it is a real detector-level performance paper for **GN2**, explicitly described here as a **transformer-based model for jet flavour tagging**. The main result is not just another architecture claim from simulation: ATLAS measures the data efficiency and shows substantially better light- and charm-jet rejection than DL1d at fixed b efficiency, with uncertainties around 1% in a useful kinematic region. That makes it relevant for understanding what kinds of modern sequence/attention architectures are actually surviving contact with calibration and deployment constraints. This might be especially interesting for you, **Aritra**, because it is directly in the lane of **jet tagging with novel ML architectures**, and it also gives a concrete benchmark for what "better than the previous generation" means in Run 3 conditions. Even though it is a calibration paper rather than a method-development paper, the fact that the architecture is already validated in data is exactly what makes it worth attention. I would read this as a reality check on where frontier jet taggers currently stand experimentally.

Tracking, triggers, and online reconstruction with ML

- **Hybrid pattern recognition for charged particle tracking: Hough transform and convolutional neural efficiency networks**
  This paper combines a geometry-driven **Hough transform** front end with a neural network that suppresses false track candidates in HL-LHC-like occupancy. The novelty is the division of labor: the Hough stage keeps the fast, detector-structured candidate generation, while the learned stage operates on the Hough parameter-space image rather than on a heavily re-encoded representation. That is a sensible design for online or near-online reconstruction, because it targets the actual bottleneck of spurious candidates rather than replacing the whole chain with a black box. This might be interesting for you, **Aritra**, as a broader example of ML integrated into collider reconstruction in a way that respects detector geometry and latency constraints; **Jingjing**, you might also care because this is the kind of analysis-workflow component where future automation and agentic tooling could plug into reconstruction optimization. It is not one of your highest-priority topics, but it is one of the more technically substantive ML-for-HEP papers in today's list.

- **Neural-Network-Assisted Binary Template Construction for Matrix-Based Pattern Matching in the STCF MDC**
  The useful idea here is that the neural network is used **offline to optimize binary template libraries**, but there is **no neural-network inference at runtime**. That is a clever compromise for trigger or HLT environments: you get learned optimization of the pattern bank while preserving deterministic, bitwise, parallelizable online execution. The paper frames template construction as a differentiable multi-objective optimization problem, which is more interesting than a generic classifier because it directly targets the design of the deployed algorithm. This might be interesting for you, **Jingjing**, because it is a concrete example of AI assisting an analysis or reconstruction workflow rather than replacing it end-to-end; **Aritra**, you may also find it useful as a model for how ML can be inserted into collider pipelines under strict operational constraints. I would not rank it above the GN2 paper for your interests, but it is one of the better examples today of practical AI-enhanced HEP systems design.

Simulation-based inference and posterior learning

- **Identifying lensed gravitational waves with physics-informed posterior learning**
  This is outside collider physics, but methodologically it is one of the strongest matches for **simulation-based inference** interests. The authors train an approximate posterior for shared source parameters and then **fuse posterior samples with direct waveform features** to rank lensed multi-image events against unrelated overlapping mergers. What is novel is the explicit use of a **physics-informed consistency structure**: lensing can alter amplitudes, delays, and Morse phases while preserving intrinsic phase evolution, and the ML system is built around that constraint rather than around generic pattern recognition. This might be especially interesting for you, **Jingjing**, because it is a clean example of posterior-learning machinery being used in a realistic scientific discrimination task, with held-out lens-family evaluation and performance quoted at fixed false-positive rate. It is not particle physics, so I would not put it at the very top of your queue, but if you are tracking where SBI-style methods are going in fundamental science, this is worth a look.

Di-Higgs and Higgs self-coupling searches

- **Search for nonresonant triple Higgs boson production in the final state with six bottom quarks in proton-proton collisions at $\sqrt{s}$ = 13 TeV**
  This is not a di-Higgs paper, but it is close enough to your Higgs-self-coupling interests that I would still flag it. The main value is phenomenological and experimental: CMS sets the strongest current constraint on nonresonant **HHH** production and, importantly, excludes part of the (κ_3,κ_4) space allowed by perturbative unitarity. For you, **Jingjing**, this matters because it sharpens the broader Higgs self-coupling landscape that sits adjacent to **HH searches and sensitivity studies**, even though the abstract does not advertise a specific ML contribution. I would treat this as a context paper rather than a methods paper, but it is still likely relevant if you are watching how multi-Higgs constraints are evolving experimentally. If time is short, this is probably a skim for results and limits rather than a deep read.

Symbolic regression and interpretable ML in fundamental science

- **Inverse-k Primordial Oscillations from a Symbolic Regression Search**
  This is another non-collider paper, but the methodological angle is notable: the authors use **symbolic regression** for a template-free search and recover an interpretable inverse-k oscillation form preferred over standard linear or logarithmic templates. The reason I am flagging it is not because it directly matches your core HEP programs, but because it is a good example of ML being used for **interpretable scientific model discovery** rather than just classification or regression. This might be mildly interesting for you, **Jingjing**, especially if you are thinking about agentic or automated scientific workflows that generate candidate functional forms or summary models. It is clearly lower priority than the posterior-learning paper above, since it is neither particle physics nor inference for collider data, but it is one of the few papers today with a genuinely distinct ML angle. I would file it under "methodological inspiration" rather than immediate relevance.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 47 papers hep-ex: 23 papers hep-th: 49 papers