

Hi all,

Here's what you folks might find interesting today.

TOP PICKS

Jet tagging and jet representation learning

- [What Do Lorentz-Equivariant Jet Taggers Learn?](#)

This is probably the clearest match of the day for you, Aritra. The paper does not just present another jet tagger, it opens up the black box of **Lorentz-equivariant architectures** by using equivariance tests, linear probes, and grade ablations across models like **L-GATr** and **LLoCa-T**. The key physics takeaway is that these models appear to actively suppress frame-dependent pseudorapidity information while strongly encoding more meaningful jet observables such as **jet mass** and **N-subjettiness**. The grade-ablation result is also quite concrete: for top tagging, bivector channels seem largely irrelevant, while vector-like channels carry most of the discriminative power, though in a seed-dependent way. That makes this especially relevant if you are thinking about what equivariant models are really buying us beyond performance numbers, and how to design leaner or more interpretable architectures going forward. Jingjing, this may also be worth a look for the broader question of how ML architectures can be made more physically structured and diagnostically transparent.

Agentic AI and workflow automation for particle physics

- [Towards LLM-Powered Automation of a Dark Matter Constraint Repository](#)

This feels especially relevant for you, Jingjing, because it is one of the more concrete **agentic AI for HEP workflows** papers in today's list. The authors build an end-to-end pipeline that monitors arXiv, extracts exclusion curves, converts them into code, and opens pull requests for human review, so this is much closer to analysis infrastructure automation than generic LLM commentary. What stands out technically is the use of **consensus voting over noisy extractions** plus a physics-convention canonicalization layer, which is exactly the sort of detail that determines whether these systems are useful in practice. The benchmark numbers are also nontrivial: 90.5% coupling-type classification accuracy and a median residual of 0.33 dex on a 346-paper test set. I also like that the abstract is honest about the unresolved governance problem: the pipeline is deployed and has generated proposals, but none have merged yet. That makes it useful not only as a methods paper, but as a case study in where agentic systems are already close to helping particle-physics curation and where the bottlenecks remain.

Simulation-based inference and likelihood-free methods

- [Nonparametric Deconvolution and Denoising using Simulation Based Inference](#)

This one is not particle-physics specific, but it is a strong **general-methods** paper that could matter for you, Jingjing, given your interest in **simulation-based inference**. The setup is likelihood-free deconvolution under measurement error, where the authors train a latent generative model by matching the observed distribution to a noise-convolved model distribution using a **convolutional MMD loss**. What makes it more than a generic ML paper is that they also provide finite-sample theory and convergence rates, including the expected distinction between ordinary-smooth and super-smooth inverse problems. In practical terms, the framework supports expressive latent models like **Gaussian mixtures and normalizing flows**, and then uses the learned density as an empirical prior for denoising latent variables. That combination of flexible generative modeling plus explicit inverse-problem guarantees is exactly the kind of thing that could transfer to unbinned or detector-smearred inference problems in HEP. Aritra, this may also be worth bookmarking as a broader statistics paper with possible downstream use in reconstruction or weakly supervised settings.

- [**Action-BED: Task-Driven Bayesian Experimental Design with Singly Intractable Objectives**](#)

This is another methods-heavy paper that seems especially aligned with your interests, Jingjing. The main conceptual move is to reformulate Bayesian experimental design around **expected future loss on downstream actions**, rather than the usual information-gain objective, and in doing so convert a doubly intractable problem into a **singly intractable** one. That matters because it potentially makes design optimization much more practical in simulation-based settings where explicit posterior or marginal-likelihood estimation is painful. The framework is also naturally implicit: you only need to sample from the joint model and evaluate the downstream loss, which is a very attractive property for complex scientific simulators. For particle physics, this could be relevant wherever one wants to optimize data-taking, summary choices, or analysis strategy directly for a target decision rather than for generic posterior sharpness. It is not HEP-specific, but it is exactly the sort of foundational inference machinery that could become useful in high-dimensional SBI workflows.

Reinterpretation and distributable likelihoods

- [**Accelerating Discovery: Model-Agnostic Likelihoods for the Reinterpretation of Particle Physics Results and their Application to the Belle II \$B^+ \rightarrow \bar{K}^+ \nu \bar{\nu}\$ Measurement**](#)

This looks quite relevant for you, Jingjing, even though it is framed as reinterpretation rather than SBI per se. The thesis develops a way to produce **model-agnostic likelihoods** by reweighting the joint distribution of reconstruction and kinematic variables using theory-ratio information, with the goal of preserving correlations and systematics while avoiding full reanalyses. That is a useful middle ground between expensive experimental reruns and oversimplified reinterpretation products. The concrete application is Belle II's **$B^+ \rightarrow \bar{K}^+ \nu \bar{\nu}$** measurement, but the broader methodological point is the public release of distributable likelihood objects that can support rapid inference over alternative parameter points. If this works robustly, it is exactly the sort of infrastructure that could make high-dimensional reinterpretation and likelihood-free workflows much more scalable. Even though it is not one of your headline topics, Aritra, it may also be worth a skim as an example of how to package analyses for broader downstream reuse.

Di-Higgs and multi-Higgs searches

- [**Searches for resonant and nonresonant multi-Higgs production in ATLAS and CMS**](#)

This is the clearest direct **di-Higgs** item in today's set, so I wanted to flag it for you, Jingjing. It is a broad experimental overview rather than a new ML-driven analysis, covering both resonant and nonresonant Higgs-pair production as well as triple-Higgs channels across ATLAS and CMS. The value here is mainly as a compact status update on where the field stands in constraining the Higgs self-couplings and related BSM scalar scenarios. Since your priority is especially on HH studies with a clear ML component, this is not as strong a match as it would be if it discussed specific analysis machinery, but it is still squarely on-topic phenomenologically. If you are tracking the evolving search landscape or looking for context for future ML-enabled HH work, this is likely worth a quick read. I would treat it as a situational-awareness paper rather than a methods paper.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 46 papers hep-ex: 21 papers stat.ML: 40 papers