Hi all,

Here's what you folks might find interesting today.

Anomaly detection, simulation-based inference, and unbinned likelihoods

- **Factorizable Normalizing Flows for parameter-dependent density morphing**
  This paper introduces **Factorizable Normalizing Flows**, a way to model how a density changes with continuous nuisance or physics parameters without needing training samples on the full combinatorial grid of parameter settings. The key idea is to learn each parameter's deformation separately and then compose them additively at inference time, which makes the method scale linearly rather than exponentially in the number of parameters. That is especially relevant for **unbinned high-dimensional analyses** and likelihood-free workflows where systematic morphing is the real bottleneck. This might be especially interesting for you, **Jingjing**, because it is directly aimed at the next generation of **unbinned inference** in HEP and keeps the likelihood tractable. **Aritra**, this could also be worth a look if you are thinking about model-agnostic searches where robust background density modeling under systematics matters.

- **Defining a Minimum Resolution for Unbinned Analyses**
  Szewc proposes a **Minimum Resolution Likelihood (MRL)** construction that turns hard-to-quantify systematic biases from ML-based background likelihood estimation into something more like statistical uncertainty by defining a fiducial signal region. The practical point is not just "better inference", but a concrete prescription for when unbinned ML-driven likelihood estimates stop being trustworthy at too-fine resolution. What makes it particularly relevant is that the paper includes a **realistic HI-SIGMA application to di-Higgs searches**, so it connects methodology to a target analysis rather than staying at toy level. This seems very much in your lane, **Jingjing**, both for **simulation-based / unbinned inference** and for the explicit **di-Higgs** connection. It may also be useful for you, **Aritra**, as a statistical safeguard for ML-heavy search strategies where background modeling systematics can otherwise masquerade as anomalies.

- **Propagating data noise through the fit: the Monte Carlo replica distribution**
  This is a more statistics-heavy paper, but the core result is clean: it derives exactly how the **Monte Carlo replica method** departs from the Bayesian posterior once the fitted model is nonlinear. The difference is traced to a single computable matrix, the residual-weighted Hessian at the best fit, which gives a concrete diagnostic for when replica-based uncertainties are over- or under-estimated. That makes it potentially useful well beyond PDFs and SMEFT, because many ML-assisted inference pipelines in HEP rely on approximations whose uncertainty interpretation is murky. This might be worth your attention, **Jingjing**, as background reading for inference methodology, especially if you care about when approximate uncertainty propagation is actually faithful. It is less directly on your priority list than the two papers above, but it has real methodological value.

Di-Higgs and Higgs-pair phenomenology

- **Analytic two-loop electroweak corrections at high energies**
  This proceedings-style paper summarizes recent progress on **analytic two-loop electroweak amplitudes** in the full SM, with **Higgs boson pair production** as the showcase application. The interesting bit is the emphasis on the high-energy regime, where logarithmic and power corrections develop nontrivial structure and can become sizeable, which matters for precision HH phenomenology and future collider extrapolations. It is not an ML paper, so it sits below the inference papers for your profile, but the HH focus makes it relevant. This might be useful for you, **Jingjing**, if you want to keep tabs on theory systematics that could feed into projected sensitivity studies or interpretation work around di-Higgs. I would treat it as a targeted phenomenology read rather than a must-read.

- **Influence of the KK graviton decay into hh on the triple Higgs measurement at LHC**
  This one is about interpreting possible excess structure in terms of **KK gravitons decaying**

to hh, with implications for **triple Higgs coupling measurements** and related heavy resonance searches. The relevance here is mainly phenomenological: if you are tracking the broader HH/triple-Higgs landscape, it is useful as a speculative interpretation paper connecting excesses, resonance structure, and future Run 3 prospects. There is no ML component in the abstract, so it does not score as highly as the inference-focused entries. Still, this may be of some interest to you, **Jingjing**, because it touches the HH ecosystem and possible beyond-SM contamination of Higgs self-coupling measurements. I would file it under "situational awareness" rather than priority reading.

Detector and reconstruction ML for collider experiments

- **Detector-aware target definitions for full-event particle reconstruction**
This paper tackles a very practical but often underappreciated issue in end-to-end reconstruction: the training targets themselves can be unphysical if they ignore detector resolution and geometry. The authors define **detector-aware targets** for hit-level particle-flow reconstruction, including a PF-aware merging scheme, and show that this improves not just nominal performance but also **robustness under changes in event topology**. That robustness angle is the most interesting part, because it gets at distribution shift and process dependence rather than just squeezing out a better benchmark number. This might be interesting for you, **Aritra**, since it is a GNN-based collider ML paper with a strong emphasis on physically meaningful representations and generalization. **Jingjing**, you might also care if you are thinking about analysis pipelines where reconstruction-level inductive biases feed into downstream inference.

General ML methods with possible HEP relevance

- **Calabi-Yau Metrics with Full Moduli Dependence**
This is not particle-phenomenology ML, but it does combine **machine-learned numerical data with symbolic regression** to recover approximate analytic expressions with explicit moduli dependence. The methodological novelty is the hybrid pipeline: use ML to get accurate numerical structure, then use symbolic regression to compress that into interpretable formulas. That kind of ML-to-analytic distillation is broadly interesting and could inspire analogous workflows in HEP where one wants interpretable surrogates rather than black-box emulators. This is probably more of a **general methods** paper than a direct fit for either of your core topics, but **Jingjing** may find the surrogate-modeling angle useful, and **Aritra** might appreciate the emphasis on extracting structure rather than only prediction. Not a top-priority read, but a plausible cross-pollination paper.

Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 51 papers hep-ex: 15 papers hep-th: 59 papers