I'll evaluate the papers against the researcher profiles before writing the digest.

Paper 1 (hep-ph [1]): "Local Conformal Predictions for Calibrated Surrogates" - This is about ML uncertainty quantification for LHC surrogates (scattering amplitudes/event generation). Tilman Plehn is a starred author. This maps well to Jingjing's interest in ML uncertainty quantification and unfolding, and tangentially to simulation-based inference. Score: ~7 for Jingjing. Not strongly relevant to Aritra's topics.

Paper 2 (hep-ex [1]): TMD measurements - purely phenomenological/experimental QCD, no ML component, not related to either researcher's topics. Score: ~0.5.

Paper 3 (hep-th [1]): Chebyshev approximations for Feynman integrals - numerical methods for loop integrals, no ML, not related to either researcher's topics. Score: ~0.5.

Only Paper 1 clears the bar for inclusion. Papers 2 and 3 are not relevant enough for the digest.

Hey there,

Here's what caught my eye from today's arXiv listings -- it's a light day, but there is one paper worth flagging.

ML Uncertainty Quantification

[Local Conformal Predictions for Calibrated Surrogates] *Suprio Dubey, Henning Bahl, Anja Butter, Jurgen Hesser, Tilman Plehn*

This one is for you, Jingjing. The paper tackles a genuinely hard problem: getting *locally calibrated* uncertainty estimates out of neural network surrogates for LHC scattering amplitudes, where the systematics are non-Gaussian and standard conformal prediction methods fall flat. The key contribution is FALCON, a new conformal prediction scheme that learns spatially varying confidence intervals rather than applying a single global correction -- a meaningful step beyond vanilla conformal coverage guarantees. The motivation is ultra-fast event generation, so this sits at the intersection of surrogate modelling and rigorous uncertainty quantification, both of which are directly relevant to your work on simulation-based inference and ML uncertainty quantification. It is worth noting that **Tilman Plehn** (a starred author) is on this paper, and the distribution-free nature of the approach means it could in principle be ported to other surrogate or density-estimation contexts beyond amplitude emulation. Worth a read if you are thinking about how to attach trustworthy error bars to any learned approximator in a physics pipeline.

The rest of today's listings (TMD factorisation in e+e- and SIDIS, and Chebyshev approximations for Feynman integrals) are solid theory/phenomenology work but do not connect to the ML and data-analysis topics on your radar -- skipping those.

Until the next announcement!

P.S. In the most recent submission window: hep-ph: 34 papers hep-ex: 10 papers hep-th: 38 papers