

Hello folks,

Here's what you might find interesting today. It's a relatively quiet day on the arXiv, but there's one paper that stands out.

Statistical Inference and Likelihood Methods

[The Well-Tempered Likelihood: Honest Confidence Intervals for Misspecified Models](#) -- This might be interesting for you, Jingjing. Benjamin Nachman and Jesse Thaler (both starred authors) tackle a fundamental problem in **likelihood-based inference**: when your model is misspecified, standard confidence intervals shrink to zero width as you add more data, giving overconfident and misleading results. Their "well-tempered likelihood" divides the likelihood-ratio test statistic by a goodness-of-fit measure evaluated at the best-fit point, so that under misspecification the interval self-limits at a floor set by the model's inadequacy -- effectively reducing the effective sample size. The unbinned formulation uses a **classifier two-sample test** as the GOF statistic, which directly connects to **density ratio estimation** techniques central to simulation-based inference. They demonstrate the method on a Gaussian toy example and on a synthetic strong-coupling measurement in e^+e^- collisions. The core idea -- that all models are correct as long as your dataset is small enough -- has broad implications for any **unbinned high-dimensional analysis** where model misspecification is the norm rather than the exception.

That's all worth flagging today. Back tomorrow with more!

P.S. In the most recent submission window: hep-ph: 31 papers hep-ex: 12 papers hep-th: 45 papers