Hey all,

Here's what you folks might find interesting today. It's a fairly quiet day on the arXiv for our core topics, but a few papers are worth a look.

Physics-Constrained Neural Networks and Uncertainty Quantification

- S-matrix informed neural networks for amplitude analysis by Wyatt A. Smith, Arkaitz Rodas, Marius D. Thomas, et al. This might be interesting for you, Jingjing. The authors introduce S-matrix informed neural networks (SINNs) that learn scattering amplitudes directly from data while respecting unitarity, analyticity, and crossing symmetry -- no fixed functional form required. What stands out is their novel data selection procedure: they use the response of constrained NN ensembles to identify mutually compatible experiments, effectively turning inconsistency between datasets into a selection criterion rather than a nuisance. The full pipeline unifies physics-constrained representation learning, data selection, and correlated uncertainty quantification, which directly touches your interest in ML uncertainty quantification. While the application is to pion-pion scattering rather than collider physics, the methodology -- particularly the ensemble-based uncertainty propagation and closure testing -- is transferable to other constrained inverse problems in particle physics.

- Neuro-dispersive extractions of light-meson resonances by the same author group. This is the companion paper to the SINN work above, and you might also find it interesting, Jingjing. Here they perform the first dispersive extraction of resonant poles from analytically continued neural networks, obtaining robust determinations of the sigma/f0(500), rho(770), and f0(980) poles. The key methodological point is that controlling representation dependence through the SINN framework allows them to propagate correlated uncertainties from a large network ensemble to all derived observables, including scattering lengths and Adler zeroes that emerge as predictions of the learned analytic structure. The stability against architecture variations is a useful data point for anyone thinking about how model choice affects physics-constrained ML outputs.

Graph Neural Networks for Detector Physics

- Exploring ESSnuSB Near Water Cherenkov Detector Designs Through Graph Neural Network Flavour Identification by J. Aguilar, M. Anastasopoulos, et al. This might catch your eye, Aritra. The paper applies graph neural network classifiers to neutrino flavour identification in a water Cherenkov detector, studying how classification performance degrades as detector volume and PMT coverage are reduced. The GNN formulation is natural here since detector hits can be represented as a graph, and the authors show that classification remains accurate even with volumes up to 8x smaller than nominal. While this is neutrino physics rather than LHC jets, the question of how graph-based architectures perform under varying detector granularity and instrumentation is methodologically relevant to your interest in GNN architectures for particle physics -- and the systematic study of performance versus detector design parameters is a nice template for similar studies at collider experiments.

Until the next announcement!

P.S. In the most recent submission window: hep-ph: 26 papers hep-ex: 9 papers hep-th: 38 papers