




Simulation-based inference for the stochastic gravitational-wave background

Anshu Shaw

Motivation and LISA data model


An SGWB carries information about unresolved sources and the early Universe. We infer a power-law signal together with instrumental noise from the LISA A, E and T channels [1]. The model has four parameters: amplitude, tilt, and noise levels A and P. Mock spectra follow the averaged-periodogram χ^2 distribution [2].

GNN: learned summaries



The two-dimensional summaries vary smoothly with the signal and noise parameters, indicating informative compression. [4]

Training and validation (GNN)




Training and validation losses track each other and flatten through approximately 80 epochs in the displayed run. [4]

References [1] J. Alvey et al., Phys. Rev. D 109, 083008 (2024). [2] G. Franciolini et al., arXiv:2505.24695 (2025). [3] C. Durkan et al., Neural Spline Flows, arXiv:1906.04032 (2019). [4] All figures: Shaw, Pieroni, Bertacca, Bartolo, in prep.

Neural posterior estimation


A graph neural network (GNN) treats A, E and T as nodes. Channel encoders combine logarithmic and linear spectra, then message passing and pooling produce two-dimensional summaries. Neural spline flows learn conditional posteriors jointly with the compression [1,3].

GNN: posterior comparison



Neural, MCMC and Fisher contours overlap, with residual shifts in the signal sector. Black lines mark the injected parameters. [4]

Twin-network consistency (GNN)




Noise-network disagreement decreases overall. Signal-network KL remains non-zero and fluctuates across checkpoints. [4]

Validation and scope


Compare neural posteriors with MCMC using the Whittle likelihood and with the Fisher approximation [2]. Loss curves, twin-network KL and coverage test distinct aspects of performance. The toy model assumes stationary Gaussian processes and a sky-averaged LISA response.

Physics-informed: posterior comparison



The physics-informed run provides a comparison with the GNN result for the same four-parameter signal-plus-noise model. [4]

Posterior calibration (GNN)



Coverage is close to the diagonal for most parameters. The signal amplitude shows mild overconfidence in the tails. [4]