Austrian Numerical Analysis Day 2026

Efficient sampling of deep Gaussian processes via wavelet sparsity

Conor Osborne[1], Michael Feischl[1]

[1] TU Wien

Deep Gaussian processes are hierarchical extensions of classical Gaussian processes and provide a flexible framework for modelling complex, non-stationary random fields. A key obstacle in practice, however, is the high cost of sampling, which in standard implementations scales cubically in the number of discretisation points at each layer.

We present a wavelet-based sampling strategy in which the covariance operators are represented in a multiresolution wavelet basis, yielding sparse approximations and reducing this cost to near-linear complexity.

We derive error estimates for these approximations and study how they propagate through successive layers of the model. Numerical experiments show that the method delivers substantial computational savings with controlled approximation error.