XUEMENG WANG, Simon Fraser University

Into the danger zone: stable extrapolation in high-dimensional function and operator learning

Out-of-distribution (OOD) generalization is a central challenge in scientific machine learning (SciML). We study regression problems in which the test distribution differs from the training distribution and ask: under what assumptions on the target function or operator is stable extrapolation possible, and how far beyond the training domain can one extrapolate? Existing theory often controls the test error through additive penalties measuring the discrepancy between the training and test distributions. Such guarantees can be considerably more pessimistic than the generalization observed empirically.

Under holomorphic regularity assumptions on the target, we identify classes of functions and operators for which the OOD generalization error converges at algebraic rates. This phenomenon stems from the increasing smoothness of higher-index coordinates, leading to what we term a 'blessing of high dimensionality' We derive explicit rates for arbitrary test measures supported on suitable domains and quantify how the admissible domain depends on the underlying regularity. Our extrapolation guarantees are independent of the test distribution, depending only on its support. We establish matching rates for polynomial approximation, deep neural networks, and deep neural operators. In addition, we present extensive numerical experiments across a range of functions and operators that corroborate the predicted rates.