OLIVIER ZAHM, Massachusetts Institute of Technology
Certified dimension reduction in nonlinear Bayesian inverse problems

A large-scale Bayesian inverse problem has a low effective dimension when the data are informative only on a low-dimensional subspace of the input parameter space. Detecting this subspace is essential for reducing the complexity of the inverse problem. For example, this subspace can be used to improve the performance of Markov chain Monte Carlo algorithms and to facilitate the construction of surrogates for expensive likelihood functions. Several different methods have recently been proposed to construct such a subspace—in particular, gradient-based approaches like the likelihood-informed subspace method and the active subspace method. In the context of nonlinear inverse problems, however, both methods are essentially heuristics whose approximation properties, relative to an optimal approximation, are poorly understood. We introduce a best approximation problem for the posterior distribution, which defines the best possible parameter space decomposition. Solving this problem in general non-Gaussian/non-linear cases, however, appears intractable. Instead, we develop a new bound on the Kullback-Leibler divergence between the posterior distribution and the approximation induced by a parameter space decomposition. We then identify the subspace that minimizes this bound, and compute it using gradients of the likelihood function. This approach allows the approximation error to be rigorously controlled. We also address the question of efficient computation of the parameter space decomposition when only a limited number of gradient evaluations are allowed. A numerical comparison with existing methods favorably illustrates the performance of our new approach.

This is joint work with Youssef Marzouk, Tiangang Cui, Kody Law, and Alessio Spantini.