Variational Reformulation of the Uncertainty Propagation Problem in Linear Partial Differential Equations

Ilias Bilionis, Purdue University

Stochastic partial differential equations (SPDE) have a ubiquitous presence in computational science and engineering. Stochasticity in SPDEs arises from unknown/random boundary/initial conditions or field parameters (e.g., the permeability of the ground in flow through porous media, the thermal conductivity in heat transfer) and, thus, it is inherently high-dimensional. In this regime, traditional uncertainty propagation techniques fail because they attempt to learn the high-dimensional response surfaces (the curse of dimensionality). The only viable alternative is Monte Carlo (and advanced variants such as multi-level MC). However, as A. O'Hagan put in in his seminal 1987 paper, "Monte Carlo is fundamentally unsound" because it fails to identify and exploit correlations between the samples. In this work, we develop a promising alternative to MC inspired by recent advances in probabilistic numerics (PN) and variational inference (VI). Our method does not rely on a traditional PDE solver, and it does not attempt to learn a response surface. Instead, we use PN which results in two advantages. First, we gain control over the computational cost, albeit at the expense of additional (but quantified) epistemic uncertainty. Second, PN allows us to quantify the information loss between the true solution of the uncertainty propagation problem and a candidate parameterization. The latter results in a reformulation of the uncertainty propagation problem as a variational inference problem, i.e., as a stochastic optimization problem.