Surrogate and Reduced Modeling for High-Dimensional Inverse Problems

Large-Scale Optimization for Bayesian Inference in Complex Systems

Jinglai Li, Chad Lieberman, Youssef Marzouk, Karen Willcox
Massachusetts Institute of Technology
77 Massachusetts Avenue, Cambridge, MA 02139

Abstract

This talk presents new methods in surrogate and reduced-order modeling for large-scale inverse problems. We present the theory for a new goal-oriented "inference-for-prediction" approach. Our approach is motivated by the observation that problems calling for prediction of quantities of interest via estimation of parameters are manifested by low-dimensional input–output maps from experimental data to output prediction. For many years, model reduction methodologies have been applied in the setting of forward problems where high-dimensional states evolve on a low-dimensional manifold defined by the input–output map. Our goal-oriented approach exploits a similar structure in inverse problems where the parameter is high-dimensional but the number of modes relevant to the prediction is low-dimensional.

In the case where experimental outputs and predicted outputs are linear in the parameters, we can formalize this approach by exploiting the observability gramian of control theory. For each set of outputs, there is an associated observability gramian whose dominant eigenspace reveals the modes with highest output energy. Using properties of these eigenspaces, we show that inversion can be carried out in a low-dimensional manifold at lower computational cost without sacrificing accuracy in prediction of output quantities of interest. This approach has direct theoretical analogies with the control theoretic concepts of minimal realization and balanced truncation. Our approach also extends to statistical inverse problem formulations that account for uncertainties. In the case of a Bayesian formulation of the inference problem with a Gaussian prior, Gaussian likelihood, and a linear model, we show that the posterior predictive mean is the same as the solution arising from the regularized inverse problem and that the posterior predictive covariance can be obtained by a small number of repeated solutions of the regularized inference-for-prediction problem.

In the nonlinear case, surrogate or reduced-order models of the parameter-to-observable map play an essential role in making statistical inversion computationally tractable. Yet constructing surrogates that are accurate over the entire parameter space (or in a Bayesian setting, the entire prior support) is often impractical. This is particularly the case when the support of the model parameters is large—e.g., when little prior information is available—and when the forward model is highly nonlinear on this range.

We present a new adaptive method for surrogate construction, which uses stochastic optimization methods to construct surrogates that are accurate over the support of the *posterior* distribution. Since the posterior reflects information gain relative to the prior, posterior probability mass is often concentrated in a very small region of the parameter space. High dimensionality accentuates this difference, as the ratio of suitably-defined prior and posterior hypervolumes may increase exponentially with the number of parameters. The posterior is of course not known *a priori*, and thus the construction of surrogates proceeds iteratively. We use the cross-entropy method and adaptive importance sampling to identify a sequence of biasing distributions, on which we construct localized polynomial surrogates for the parameter-to-observable map. The final biasing distribution not only determines the final surrogate; it also facilitates efficient posterior exploration via importance sampling or Markov chain Monte Carlo. The technique is *data-driven* in that the surrogate model, like the posterior distribution itself, depends on the available observations.