Journal of Computational Physics 228 (2009) 1862–1902

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems

Youssef M. Marzouk [a,*], Habib N. Najm [b]

[a] Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[b] Sandia National Laboratories, Livermore, CA 94551, USA

ARTICLE INFO

Article history:
Received 31 October 2007
Received in revised form 24 September 2008
Accepted 14 November 2008
Available online 7 December 2008

Keywords:
Inverse problems
Bayesian inference
Dimensionality reduction
Polynomial chaos
Markov chain Monte Carlo
Galerkin projection
Gaussian processes
Karhunen–Loève expansion
RKHS

ABSTRACT

We consider a Bayesian approach to nonlinear inverse problems in which the unknown quantity is a spatial or temporal field, endowed with a hierarchical Gaussian process prior. Computational challenges in this construction arise from the need for repeated evaluations of the forward model (e.g., in the context of Markov chain Monte Carlo) and are compounded by high dimensionality of the posterior. We address these challenges by introducing truncated Karhunen–Loève expansions, based on the prior distribution, to efficiently parameterize the unknown field and to specify a stochastic forward problem whose solution captures that of the deterministic forward model over the support of the prior. We seek a solution of this problem using Galerkin projection on a polynomial chaos basis, and use the solution to construct a reduced-dimensionality surrogate posterior density that is inexpensive to evaluate. We demonstrate the formulation on a transient diffusion equation with prescribed source terms, inferring the spatially-varying diffusivity of the medium from limited and noisy data.

1. Introduction

Inverse problems arise from indirect observations of a quantity of interest. Observations may be limited in number relative to the dimension or complexity of the model space, and the action of the forward operator may include filtering or smoothing effects. These features typically render inverse problems ill-posed—in the sense that no solution may exist, multiple solutions may exist, or solutions may not depend continuously on the data. In practical settings, where observations are inevitably corrupted by noise, this presents numerous challenges.

Classical approaches to inverse problems have used regularization methods to impose well-posedness, and solved the resulting deterministic problems by optimization or other means [1]. However, important insights and methodologies emerge by casting inverse problems in the framework of statistical inference [2,3]. Here we focus on Bayesian approaches, which provide a foundation for inference from noisy and limited data, a natural mechanism for regularization in the form of prior information, and in very general cases—e.g., nonlinear forward operators, non-Gaussian errors—a quantitative assessment of uncertainty in the results [4,5]. Indeed, the output of Bayesian inference is not a single value for the quantity of interest, but a probability distribution that summarizes all available information about this quantity, be it a vector of parameters or a function (i.e., a signal or spatial field). Exploration of this posterior distribution—and thus estimating means, higher

* Corresponding author. Tel.: +1 617 253 1337.
E-mail addresses: ymarz@mit.edu (Y.M. Marzouk), hnnajm@sandia.gov (H.N. Najm).

0021-9991/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2008.11.024

moments, and marginal densities of the inverse solution—may require repeated evaluations of the forward operator. For complex physical models and high-dimensional model spaces, this can be computationally prohibitive.

Our previous work [6] sought to accelerate the Bayesian solution of inverse problems through the use of stochastic spectral methods. Based on polynomial chaos (PC) representations of random variables and processes [7–12], stochastic spectral methods have been used extensively for *forward* uncertainty propagation—characterizing the probability distribution of the output of a model given a known distribution on the input. These methods constitute attractive alternatives to Monte Carlo simulation in numerous applications: transport in porous media [13], structural mechanics [14], thermo-fluid systems [15–17], electrochemical microfluid systems [18], and reacting flow [19]. In the *inverse* context, the Bayesian formulation in [6] constructs a stochastic forward problem whose random inputs span the support of the prior and seeks its solution using Galerkin methods. The prior support may be partitioned, but for each partition the stochastic forward problem is solved only once. The resulting spectral representations of the forward operator enter the likelihood function, and exploration of the posterior is recast as Monte Carlo sampling of the variables underlying the PC expansion. We used this scheme to infer parameters appearing nonlinearly in a transient diffusion equation, demonstrating exponential convergence to the true posterior and substantial speedup.

Other attempts at accelerating Bayesian inference in computationally intensive inverse problems have relied on reductions or surrogates for the forward model, constructed through repeated forward simulations. Wang and Zabaras [20] use proper orthogonal decomposition (POD) [21] to accelerate forward model calculations in a radiative source inversion problem. The empirical basis used for model reduction is pre-constructed using full forward problem simulations. The choice of inputs to these simulations—in particular, how closely the inputs must resemble the inverse solution—can be important [20]. Balakrishnan et al. [22] introduce a PC representation of the forward model in a groundwater transport parameter identification problem, but obtain the PC coefficients by collocation; again, this process depends on a series of "snapshots" obtained from repeated forward simulations. In the statistical literature, under the headline of "Bayesian parameter calibration", Gaussian processes have been used extensively as surrogates for complex computational models [23]. These approaches treat the forward model as a black box, and thus require careful attention to experimental design and to modeling choices that specify the mean and covariance of the surrogate Gaussian process. A different set of approaches retain the full forward model but use simplified or coarsened models to guide and improve the efficiency of Markov chain Monte Carlo (MCMC). Christen and Fox [24] use a local linear approximation of the forward model to improve the acceptance probability of proposed moves, reducing the number of times the likelihood must be evaluated with the full forward model. Higdon et al. [25] focus on the estimation of spatially distributed inputs to a complex forward model. They introduce coarsened representations of the inputs and apply a Metropolis-coupled MCMC scheme [26] in which "swap proposals" allow information from the coarse-scale formulation to influence the fine-scale chain. Efendiev et al. [27] also develop a two-stage MCMC algorithm, using a coarse-scale model based on multiscale finite volume methods to improve the acceptance rate of MCMC proposals.

This paper extends the stochastic spectral methodology of [6] to inverse problems whose solutions are unknown functions—i.e., spatial or temporal fields. In doing so, we also explore dimensionality reduction in the Bayesian formulation of inverse problems, and the dependence of dimensionality on both the prior and the data. Inverse problems involving fields are vital to applications ranging from geophysics to medical imaging. Spatial fields may correspond to inhomogeneous material properties, such as permeabilities, diffusivities, or densities, or may represent distributed source terms in transport equations.

Estimating fields rather than parameters typically increases the ill-posedness of the inverse problem, since one is recovering an infinite-dimensional object from finite amounts of data. Obtaining physically meaningful results requires the injection of additional information on the unknown field—i.e., regularization [3]. A standard Bayesian approach is to employ Gaussian process (GP) or Markov random field (MRF) priors [28,4,29]. Most studies then explore the value of the field on a finite set of grid points [30]; the dimension of the posterior is tied to the discretization of the field. This recipe presents difficulties for stochastic spectral approaches, however, as the size of a PC basis does not scale favorably with dimension [9]. Moreover, with any degree of smoothness, the value of the field at each grid point hardly represents an independent direction.

Ideally, one should employ a representation that reflects how much information is truly required to capture variation among realizations of the unknown field. To this end, we introduce a Karhunen–Loève (K–L) expansion based on the prior random process, transforming the inverse problem to inference on a truncated sequence of weights of the K–L modes. Other recent work has also employed K–L expansions in the context of statistical inverse problems. Li and Cirpka [31] emphasize the role of K–L expansions in enabling geostatistical inversion on unstructured grids. Efendiev et al. [27] use K–L expansions to parameterize the log-permeability field in their two-stage MCMC scheme, and introduce constraints among the weights in order to match known values of the permeability at selected spatial locations. In contrast to [31], we use a fully Bayesian approach, generating true conditional realizations from a non-Gaussian posterior.

A more fundamental distinction of the present work is that we combine a K–L representation of the unknown field with spectral methods for uncertainty propagation. In particular, the Karhunen–Loève representation of a scaled Gaussian process prior *defines* the uncertainty that is propagated through the forward model with a stochastic Galerkin scheme. The deterministic forward model, originally specified by (a system of) partial differential equations, is thus replaced by stochastic PDEs; numerical approaches to such systems, in which random fields appear as boundary conditions or coefficients, have seen extensive development [9,16,32–35]. Uncertainty propagation yields a polynomial approximation of the forward operator over the support of the prior. This approximation then enters a reduced-dimensionality surrogate posterior, which we