Journal of Computational Physics 329 (2017) 91–125

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Multimodal, high-dimensional, model-based, Bayesian inverse problems with applications in biomechanics

I.M. Franck, P.S. Koutsourelakis *

Professur für Kontinuumsmechanik, Technische Universität München, Boltzmannstrasse 15, 85747 Garching (b. München), Germany

ARTICLE INFO

Article history:
Received 24 July 2016
Received in revised form 10 October 2016
Accepted 14 October 2016
Available online 20 October 2016

Keywords:
Uncertainty quantification
Variational inference
Bayesian computation
Multimodality
Inverse problems
Dimensionality reduction

ABSTRACT

This paper is concerned with the numerical solution of model-based, Bayesian inverse problems. We are particularly interested in cases where the cost of each likelihood evaluation (forward-model call) is expensive and the number of unknown (latent) variables is high. This is the setting in many problems in computational physics where forward models with nonlinear PDEs are used and the parameters to be calibrated involve spatio-temporarily varying coefficients, which upon discretization give rise to a high-dimensional vector of unknowns.

One of the consequences of the well-documented ill-posedness of inverse problems is the possibility of multiple solutions. While such information is contained in the posterior density in Bayesian formulations, the discovery of a single mode, let alone multiple, poses a formidable computational task. The goal of the present paper is two-fold. On one hand, we propose approximate, adaptive inference strategies using mixture densities to capture multi-modal posteriors. On the other, we extend our work in [1] with regard to effective dimensionality reduction techniques that reveal low-dimensional subspaces where the posterior variance is mostly concentrated. We validate the proposed model by employing Importance Sampling which confirms that the bias introduced is small and can be efficiently corrected if the analyst wishes to do so. We demonstrate the performance of the proposed strategy in nonlinear elastography where the identification of the mechanical properties of biological materials can inform non-invasive, medical diagnosis. The discovery of multiple modes (solutions) in such problems is critical in achieving the diagnostic objectives.

1. Introduction

Model-based (or model-constrained) inverse problems appear in many scientific fields and their solution represents a fundamental challenge in the context of model calibration and system identification [2]. Bayesian formulations offer a rigorous setting for their solution as they account for various sources of uncertainty that is unavoidably present in these problems. Furthermore, they possess a great advantage over deterministic alternatives, as apart from point-estimates, they provide quantitative metrics of the uncertainty in the unknowns encapsulated in the *posterior distribution* [3].

* Corresponding author.
E-mail addresses: franck@tum.de (I.M. Franck), p.s.koutsourelakis@tum.de (P.S. Koutsourelakis).
URL: http://www.contmech.mw.tum.de (P.S. Koutsourelakis).

http://dx.doi.org/10.1016/j.jcp.2016.10.039
0021-9991/

An application of particular, but not exclusive, interest for this paper involves the identification of the mechanical properties of biological materials, in the context non-invasive medical diagnosis (elastography). While in certain cases mechanical properties can also be measured directly by excising multiple tissue samples, non-invasive procedures offer obvious advantages in terms of ease, cost and reducing risk of complications to the patient. Rather than X-ray techniques which capture variations in density, the identification of stiffness, or mechanical properties in general, can potentially lead to earlier and more accurate diagnosis [4,5], provide valuable insights that differentiate between modalities of the same pathology [6], monitor the progress of treatments and ultimately lead to patient-specific treatment strategies.

All elastographic techniques consist of the following three basic steps [7]: **1)** excite the tissue using a (quasi-)static, harmonic or transient source, **2)** (indirectly) measure tissue deformation (e.g., displacements, velocities) using an imaging technique such as ultrasound [8], magnetic resonance [9] or optical tomography [10], and **3)** infer the mechanical properties from this data using a suitable continuum mechanical model of the tissue's deformation. Perhaps the most practical of existing imaging techniques, due to its lower relative cost and increased portability, is ultrasound elasticity imaging [11,12]. The pioneering work of Ophir and coworkers [8] followed by several clinical studies [13–16] have demonstrated that the resulting strain images typically improve the diagnostic accuracy over ultrasound alone. Apart from breast cancer, there is a wealth of evidence indicating the potential of elastography-based techniques in detecting a variety of other pathologies such as prostate [17,18] and liver cancer [19], characterizing blood clots [20], brain imaging [21], atherosclerosis [22] and osteopenia [23]. As the rate of data acquisition increases and the cost decreases, it becomes increasingly important to develop tools that leverage the capabilities of physics-based models in order to produce quickly and accurately diagnostic estimates as well as quantify the confidence in them.

In this paper we advocate a probabilistic, *indirect or iterative* procedure (in contrast to *direct elastography* [24]) which admits an inverse problem formulation and involves the discrepancy between observed and model-predicted displacements [25–28,7]. Several other problems which involve complex forward models (i.e., expensive likelihoods) and *spatially varying*, unknown, model parameters share similar characteristics such as permeability estimation for soil transport processes that can assist in the detection of contaminants, oil exploration and carbon sequestration [29–31].

The solution of such model calibration problems in the Bayesian framework is hampered by two main difficulties. The first affects the computational efficiency of such methods and stems from the poor scaling of traditional Bayesian inference tools, with respect to the dimensionality of the unknown parameter vector – another instance of the *curse-of-dimensionality*. In problems such as the one described above, the model parameters of interest (i.e., material properties) exhibit spatial variability which requires fine discretizations in order to be captured. This variability can also span different scales [32,33]. Standard Markov Chain Monte Carlo (MCMC, [34]) techniques require an exorbitant number of likelihood evaluations (i.e., solutions of the forward model) in order to converge [35–38]. As each of these calls implies the solution of very large systems of (non)linear, and potentially transient, equations, it is generally of interest to minimize their number particularly in time-sensitive applications. Advanced sampling schemes, involving adaptive MCMC [39–41] and Sequential Monte Carlo (SMC, [42,33,43]) exploit the physical insight and the use of multi-fidelity solvers in order to expedite the inference process. Nevertheless, the number of forward calls can still be in the order of tens of thousands if not much more. More recent treatments attempt to exploit the lower-dimensional structure of the target posterior by identifying subspaces where either most of the probability mass is contained [1] or where maximal sensitivity is observed [44–47]. This enables inference tasks that are carried out on spaces of significantly reduced dimension and are not hampered by the aforementioned difficulties. Generally, all such schemes construct approximations around the MAP point by employing local information (e.g., gradients) and are therefore not suitable for multi-modal or highly non-Gaussian posteriors.

The latter represents the second challenge that we attempt to address in this paper. That is, the identification of multiple posterior modes. In the context of elastography, multi-modality can originate from anisotropic materials [48], wrong/missing information from images/measurements [49] or the imaging modality employed [50]. In all cases, each mode in the posterior can lead to different diagnostic conclusions and it is therefore very important to identify them and correctly assess their posterior probabilities. The majority of Bayesian strategies for the solution of computationally intensive inverse problems operates under the assumption of a unimodal posterior or focuses on the approximation of a single mode of the posterior. Some numerical inference tools based on SMC or other tempering mechanisms [51–53] have been developed but require a very large number of forward model calls particularly when the dimension of unknowns increases.

In this paper we propose a Variational Bayesian (VB) strategy that extends our previous work [1]. Therein we have shown how accurate approximations of the true posterior can be attained by identifying a low-dimensional subspace where posterior uncertainty is concentrated. This has led to computational schemes that require only a few tens of forward model runs in the problems investigated. Nevertheless, our previous work was based on the assumption of a unimodal posterior which we propose overcoming in this paper by employing a mixture of multivariate Gaussians. We note that a different VB strategy that also makes use of mixtures of Gaussians to solve model-based inverse problems has been proposed in [54]. Mixture models have also been employed in various statistics and machine learning applications (e.g., speaker identification [55], data clustering [56]) in combination with Variational Bayesian inference techniques [57–59]. Nevertheless, all these problems were characterized by inexpensive likelihoods, relatively low dimensions and multiple data/measurements. In contrast, the inverse problems considered here are based on a single experiment and a single observation vector.

In addition, we propose an adaptive algorithm based on information-theoretic criteria for the identification of the number of the required mixture components (Section 2). We present the parametrization of the proposed model in Section 2 where