Journal of the Mechanics and Physics of Solids 118 (2018) 74–97

Contents lists available at ScienceDirect

Journal of the Mechanics and Physics of Solids

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

Bayesian inference of the spatial distributions of material properties

A. Vigliotti[a,b], G. Csányi[a], V.S. Deshpande[a,*]

[a] *Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK*
[b] *Innovative Materials Laboratory, Italian Aerospace Research Centre, Capua 81043, Italy*

ARTICLE INFO

Article history:
Received 22 October 2017
Revised 5 April 2018
Accepted 9 May 2018

Keywords:
Occam's razor
Inverse elasticity problems
Model selection
Nested Sampling

ABSTRACT

The inverse problem of estimating the spatial distributions of elastic material properties from noisy strain measurements is ill-posed. However, it is still typically treated as an optimisation problem to maximise a likelihood function that measures the agreement between the measured and theoretically predicted strains. Here we propose an alternative approach employing Bayesian inference with Nested Sampling used to explore parameter space and compute Bayesian evidence. This approach not only aids in identifying the basis function set (referred to here as a model) that best describes the spatial material property distribution but also allows us to estimate the uncertainty in the predictions. Increasingly complex models with more parameters generate very high likelihood solutions and thus are favoured by a maximum likelihood approach. However, these models give poor predictions of the material property distributions with a large associated uncertainty as they overfit the noisy data. On the other hand, Bayes' factor peaks for a relatively simple model and indicates that this model is most appropriate even though its likelihood is comparatively low. Intriguingly, even for the appropriate model that has a unique maximum likelihood solution, the measurement noise is amplified to give large errors in the predictions of the maximum likelihood solution. By contrast, the mean of the posterior probability distribution reduces the effect of noise in the data and predicts the material properties with significantly higher fidelity. Simpler model selection criteria such as the Bayesian information criterion are shown to fail due to the non-Gaussian nature of the posterior distribution of the parameters. This makes accurate evaluation of the posterior distribution and the associated Bayesian evidence integral (by Nested Sampling or other means) imperative for this class of problems. The output of the Nested Sampling algorithm is also used to construct likelihood landscapes. These landscapes show the existence of multiple likelihood maxima when there is paucity of data and/or for overly complex models. They thus graphically illustrate the pitfalls in using optimisation methods to search for maximum likelihood solutions in such inverse problems.

1. Introduction

Estimating the spatial distributions of mechanical properties in a heterogeneous solid body from the measurements of strains or displacements fields has wide ranging applications, including material characterisation, medical diagnosis and civil

* Corresponding author.
E-mail address: vsd@eng.cam.ac.uk (V.S. Deshpande).

https://doi.org/10.1016/j.jmps.2018.05.007

infrastructure monitoring. For example, there has been a recent burst of activity in the development of smart civil infrastructure and this includes application of strain measurement technologies like fibre optic sensing to monitor the structural health of tunnels (Gue et al., 2015), bridges (Ko and Ni, 2005) and concrete sleepers (Butler et al., 2017), to name a few. In these applications strains are measured at a small number of discrete locations. In some other applications full field measurements of the displacement fields within specimens are available. For example, a technique known as "displacements under applied loading by Magnetic Resonance Imaging" (dualMRI) has been developed to perform in vivo measurements of displacements and strain in musculoskeletal tissues (Chan et al., 2012). In all these cases the spatial distributions of material properties is the information of primary interest but displacement and strain information does not directly describe these material property distributions. The reconstruction of material property maps from noisy (and sometimes sparsely spaced) strain measurements is an ill-posed inverse problem that requires complex modelling approaches.

A number of methods have been proposed to identify constitutive parameters from strain/displacement measurements; readers are referred to Avril et al. (2008) for an overview. Two commonly used approaches are: (i) the finite element model updating (FEMU) approach and (ii) the virtual field method (VFM). The former strategy is an optimisation method and involves adjusting the parameters in order to minimise the difference between computed and measured strains as measured by a likelihood function (Rouger et al., 1991; Molimard et al., 2005). This approach can be used for either full field or discrete strain data. On the other hand, VFM (Grédiac, 1989) is a direct identification method that does not require any model iteration and hence is computationally less expensive but is best suited for full field data. Moreover, in practice, it requires a well-chosen mechanical test that excites of all strain components under the known boundary conditions. Examples of proposed specimens include a T-shaped specimen under tensile loading (Grédiac and Pierron, 1998) and thick laminated composite tubes under compression (Pierron et al., 2000). Nevertheless, there is a strong similarity between these two identification methods, at-least in the context of linear elasticity, as discussed by Avril and Pierron (2007).

The FEMU and VFM strategies discussed above are best suited for constitutive parameter identification problems in homogenous media. In order to use these methods to determine spatial property distributions (for example in heterogeneous materials such as the articular cartilage or a structurally deteriorating bridge structure), the unknown constitutive parameters need to be cast as weights of basis functions used to describe the spatial variations of the material properties. This is in fact what is done in the so-called "equilibrium gap" method and is equivalent to VFM with piecewise fields. The resulting inverse problem of identifying these weights is typically ill-posed with the outcome not only depending on the solution strategy, but also on the choice of the basis functions. Various types of regularisations have been proposed to reduce the intrinsic instability of these solutions, but the reconstructed results are inevitably strongly dependent on the choice of the regularisation parameters (Richards et al., 2009). Moreover, the instabilities are aggravated by the presence of measurement noise. Adjoint-weighted and gradient-based variational methods have recently been proposed in an attempt to stabilise the solutions in the presence of measurement noise; see for example Bal and Uhlmann (2013) and Bal et al. (2014). By contrast, Bellis (2017) proposed a reconstruction method based on an integral formulation of the linear elasticity problem whereby a given strain field is expressed as the solution of the Lippmann–Schwinger equation. This approach circumvents the underlying instability issues but is necessarily restricted to linear elasticity and elastic moduli distributions with a small contrast. By contrast, Nguyen et al. (2015) proposed a multiscale statistical inverse method for performing the experimental identification of the elastic properties of materials at macro and mesoscales. Their method allows for identification of both the mean component and the statistical fluctuations of a stochastic model of the elasticity field of a heterogeneous microstructure using experimental data from a single specimen

Statistical and probabilistic methods are in fact being increasingly used for solutions of ill-posed inverse problems. In particular, the Bayesian approach allows for a full characterisation of all possible solutions, and their relative probabilities, whilst simultaneously addressing the problems associated with the ill-posed problem in a clear and precise fashion. Readers are referred to Stuart (2010) and also the monographs by Kiapio et al. (2005) and Tarantola (2005) for a detailed discussion of the mathematical basis of Bayesian inference methods. Bayesian approaches, although computationally expensive to implement, are starting to lie within the range of the available computational resources especially given that they allow for the quantification of uncertainty in inverse problems. Most of current Bayesian approaches use Markov Chain Monte Carlo methods and/or filtering to identify maximum posterior solutions with the uncertainty quantified by evaluating the Hessian. Examples in solid mechanics include the works of Bui-Thanh et al. (2013) to derive the material property distributions via Monte Carlo sampling from seismic measurements and that of Thurin et al. (2017) who used Kalman filtering in a two-dimensional Marmousi model. The complete Bayesian calculation of evaluating the posterior distribution and computing the Bayesian evidence is avoided in these studies, presumably due to the numerical difficulties associated with such a task. Consequently, the fidelity of the implicit approximations in these Bayesian approaches remains unclear for problems in solid mechanics.

In this study we follow a Bayesian approach to the inverse problem of determining spatial material property distributions from strain measurements. In particular we do not use Laplacian or other such approximations for the posterior as in the majority of previous studies. Rather we propose a method to estimate the entire posterior probability distribution of the material properties and thereby quantify the Bayesian evidence in support of particular choices of basis functions. The outline of the paper is as follows. We first present an overview of the Bayesian inference technique as applied to the inverse problem. Next, we describe the inverse elasticity problem and the Nested Sampling technique used to evaluate posterior probability distributions and Bayesian evidence for this elasticity problem. Finally, we discuss the computational results by comparing solutions based on Bayesian inference and maximal likelihood.