Bayesian sensitivity analysis of a nonlinear finite element model

<blockquote>

W. Becker [a,*], J.E. Oakley [b], C. Surace [c], P. Gili [d], J. Rowson [a], K. Worden [a]

[a] Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK
[b] School of Mathematics and Statistics, University of Sheffield, Hicks Building, Sheffield S3 7RH, UK
[c] Department of Structural and Geotechnical Engineering, corso Duca degli Abruzzi 24, Politecnico di Torino, 10129 Torino, Italy
[d] Department of Aeronautical and Space Engineering, corso Duca degli Abruzzi 24, Politecnico di Torino, 10129 Torino, Italy

</blockquote>

ABSTRACT

A major problem in uncertainty and sensitivity analysis is that the computational cost of propagating probabilistic uncertainty through large nonlinear models can be prohibitive when using conventional methods (such as Monte Carlo methods). A powerful solution to this problem is to use an emulator, which is a mathematical representation of the model built from a small set of model runs at specified points in input space. Such emulators are massively cheaper to run and can be used to mimic the "true" model, with the result that uncertainty analysis and sensitivity analysis can be performed for a greatly reduced computational cost. The work here investigates the use of an emulator known as a Gaussian process (GP), which is an advanced probabilistic form of regression. The GP is particularly suited to uncertainty analysis since it is able to emulate a wide class of models, and accounts for its own emulation uncertainty. Additionally, uncertainty and sensitivity measures can be estimated analytically, given certain assumptions. The GP approach is explained in detail here, and a case study of a finite element model of an airship is used to demonstrate the method. It is concluded that the GP is a very attractive way of performing uncertainty and sensitivity analysis on large models, provided that the dimensionality is not too high.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

Uncertainty analysis (UA) is a field of increasing interest in computer modelling, both within and beyond the realm of engineering. Advances in the theory of numerical simulation (such as the continuing advances in finite element (FE) analysis and computational fluid dynamics), coupled with a steady increase in available processing power, have enabled the use of increasingly sophisticated simulations to model complicated real-world processes. However, the increased complexity of these models tends to require a greater amount of information to be specified in the input. Examples within the context of engineering models could include material properties, loads, dimensions and temperatures. Any of these inputs can be subject to some uncertainty; for example, the operating temperature of a structure could be within a wide range, or the material properties could vary naturally (a good example of this is in modelling biomaterials — see [1]). The question, central to UA, then arises: what is the uncertainty in the model output, given the input uncertainty?

It is often agreed that this question poses three separate issues. The first is that of *quantification*, which deals with how to express the uncertainty in the inputs mathematically. There are in fact many ways of approaching this (possibility theory, fuzzy sets and Dempster–Shafer theory to name but a few), the majority of which are summarised in a recent book by Klir [2]. A discussion of the merits of each is outside of the scope of this paper, although without doubt the most well-established framework is probability theory, which will be used hereafter. Probability theory is an excellent approach provided that probability distributions can be defined for each input, a process known as *elicitation* (the probability-specific problem of quantification). This is in itself an area of significant research — see for example [3] for an extensive discussion. The problem of elicitation will not be addressed here, however. A second issue for UA, known as *fusion*, seeks to "fuse" information from (equivalently translating between) different uncertainty frameworks — see for example Zadeh's recent paper on a "Generalised theory of uncertainty" [4], or Ross's work on "Total Uncertainty" [5]. The final issue for UA, assuming that the uncertainty is correctly quantified, concerns an obvious step, known as *propagation*, which seeks to quantify the probability distributions (and hence the uncertainty) of the model outputs. It is the propagation problem that will be addressed in detail here.

A natural progression of UA is to seek to reduce the uncertainty in the model output (equivalent to increasing the model robustness). For any given model, it is often the case that a small set of the model inputs are causing the majority of the output uncertainty (a heuristic known as the Pareto principle [6]) — to put it another way, even if equal uncertainty were assigned to all model inputs, it is typical that the output is considerably more sensitive to uncertainty in a few inputs, and practically insensitive to the remainder. Therefore, a logical step towards reducing model uncertainty is to identify which of the inputs are causing the most uncertainty, a process known as sensitivity analysis (SA). Saltelli [7] divides SA approaches into three categories. At the most basic level, *screening* ranks the inputs in order of importance in affecting the output. The next level, *local SA*, analyses and quantifies the effects of varying input parameters, but only around the immediate locality of a specified point in the input space; one approach is based on the assumption of small perturbations about an operating point such that a linear expansion can be applied. However, this does not allow for full exploration of the input space unless responses are globally linear functions of the inputs, so is of limited use in complex models. The most informative analysis is *global SA*, which investigates and quantifies uncertainties over the complete range of input space. Unsurprisingly, this comes with the drawback of increased computational expense, but is very often necessary for large engineering models, since nonlinearities are often present, or can rarely be discounted.

Uncertainty propagation and sensitivity analysis are intrinsically linked, since a global SA requires propagation of uncertainty through a model. Many techniques have been proposed to deal with both. In the engineering literature, recent popular approaches to propagating uncertainty include the spectral stochastic FE method [8], which also been extended to provide sensitivity estimates [9]. Random matrices are also a key area of development — see the work of Soize [10]. However, both of these methods require some intervention in the model (FE) code. An alternative class of methods considers the model as a "black-box" system — in other words, the only information of interest is the value of the model outputs for a given set of input values. A conventional black-box approach is the Monte Carlo method [11], however this is computationally intensive and often impractical for large models. Some improvements can be achieved by improved sampling strategies (see e.g., [12]), but analysis of very large models is still unfeasible.

In order to reduce computational expense, a class of methods exist which involve building an emulator (equivalently a *metamodel* or *surrogate* model) which imitates the behaviour of the original model, but is considerably (computationally) cheaper to run. The emulator is trained from a small number of training data, consisting of model runs at selected values of input variables — the process is therefore a form of data modelling or machine learning. It may then be used to provide uncertainty and sensitivity estimates either by applying Monte Carlo techniques to the emulator, or even better, by analytically propagating uncertainty if the emulator is sufficiently tractable. The main issue with emulator approaches is that an emulator is required that is both efficient (i.e., requires as few training data as possible), and can emulate as wide a class of model responses as possible. For example, a polynomial function could be used as a simple emulator, but requires assumptions about the order of the function, which runs the risk of over- or under-fitting the data if the assumptions are incorrect. More sophisticated approaches in the literature range from linear models (including linear combinations of basis functions, such as polynomials, Fourier expansions and other orthogonal expansions), to radial basis functions, artificial neural networks, splines and support vector machines, all of which are discussed in [13]. Some newer approaches include "gradient boosting machines" [14] (a method based on nested "weak learners"), ACOSSO [15] (a method using multivariate splines and variable selection), "multivariate adaptive regression splines" (MARS) [16] (using an optimised combination of linear splines). A further reference on emulators can be found in [17], and comparisons in [18–21].

A relatively new emulator-based approach proposed by Oakley and O'Hagan considers the problem from a Bayesian perspective [22]. In this particular method, a Gaussian process (GP) is used as the emulator for the model. GPs are particularly well-suited since they are semi- or non-parametric (depending on the exact definition), and therefore applicable to a wide class of problems. Furthermore, they are efficient and analytically tractable, under certain assumptions. This paper aims to introduce this approach to structural dynamics as an efficient and detailed approach to UA/SA.

The following sections outline the principals behind GPs, followed by an explanation of their use as part of an efficient encompassing method for performing UA and SA at a reduced computational cost. This is then applied to a case study of a nonlinear FE model of an airship. Comments follow on the efficacy of the method and its practical implementation.