

Aspects of Uncertainty Analysis for Large Nonlinear Computational Models

K. Worden^{1,a}, W.E. Becker^{1,b}, M. Battipede^{2,c} and C. Surace^{3,d}

¹ Dynamics Research Group, Department of Mechanical Engineering,
University of Sheffield, Mappin Street, Sheffield S1 3JD, UK

² Department of Aeronautical and Space Engineering, corso Duca degli Abruzzi 24,
Politecnico di Torino, 10129 Torino, Italy

³ Department of Structural and Geotechnical Engineering, corso Duca degli Abruzzi 24,
Politecnico di Torino, 10129 Torino, Italy

^ak.worden@sheffield.ac.uk, ^bw.becker@sheffield.ac.uk, ^cmanuela.battipede@polito.it,
^dcecilia.surace@polito.it

Keywords: Nonlinear models, uncertainty analysis, Bayesian sensitivity analysis.

Abstract. This paper concerns the analysis of how uncertainty propagates through large computational models like finite element models. If a model is expensive to run, a Monte Carlo approach based on sampling over the possible model inputs will not be feasible, because the large number of model runs will be prohibitively expensive. Fortunately, an alternative to Monte Carlo is available in the form of the established Bayesian algorithm discussed here; this algorithm can provide information about uncertainty with many less model runs than Monte Carlo requires. The algorithm also provides information regarding sensitivity to the inputs i.e. the extent to which input uncertainties are responsible for output uncertainty. After describing the basic principles of the Bayesian approach, it is illustrated via two case studies: the first concerns a finite element model of a human heart valve and the second, an airship model incorporating fluid structure interaction.

Introduction

The early stages of development of a new engineering structure from design to commissioning can be exceedingly costly. One of the major drivers in this cost is the need for physical prototyping and experimental testing of the prototype to make sure it meets design requirements. It is argued that these costs could be reduced considerably by more emphasis on computer modelling in the early stages of product development – a sort of virtual prototyping. However, one of the major problems anticipated in the attempt to replace testing with modelling is that computer models will not generally provide a complete representation of the in-the-field structural behaviour. Assuming that adequate computer resources are available; the main reason for lack of model fidelity will be *model uncertainty*. In this context there are two types of uncertainty: *aleatory* and *epistemic*. Aleatory uncertainty, sometimes called ‘irreducible’ uncertainty is associated with the truly random and unknowable features of the operating environment of the structure e.g. the temperature experienced by the structure. Epistemic uncertainty, sometimes called ‘reducible’ uncertainty is associated with incomplete knowledge of the structure; information which could in principle be learned by future experimentation. If predictions are to be made on the basis of a model of a structure or system, it is critical to take account of aleatory uncertainty in order to encompass the full range of possible behaviours under random conditions. It is also important to assess the contribution of uncertain model parameters to the overall response uncertainty; in this way future experimentation can be directed towards learning more about the uncertain parameters which contribute most to uncertainty in the predictions. These two problems can be accommodated respectively by *uncertainty propagation* and *sensitivity analysis*. The current paper will discuss how the required analysis can be carried out, particularly in relation to a principled Bayesian framework. The discussion is

illustrated through the use of two case studies: the first concerning a human heart valve model and the second, an airship model incorporating fluid structure interaction.

Uncertainty Propagation in Structural Dynamics

This paper will focus on a particular aspect of uncertainty analysis which has special relevance in structural dynamics. In fact, one can argue that the analysis of uncertainty raises three main questions in dynamics. First, an appropriate *model* of uncertainty is a vital element in the design and modelling of high-value engineering structures and systems - this is the problem of *quantification*. There are numerous theoretical frameworks which allow a specification of uncertainty, like: probability theory, possibility theory, Dempster-Shafer etc., and an excellent discussion of the relations between these approaches can be found in [1]. The broader problem of quantification involves selecting the appropriate theoretical framework and then assigning a quantitative measure of uncertainty or risk. Given that there is more than one way of assigning a measure of risk, the problem arises of translating between them. This problem of normalisation is broadened into one of *fusion* by posing the question: given two assignments of uncertainty from different frameworks, how can one refine each estimate in the light of the extra information from the other? Equally important in the design process is a prescription for deciding how a measure of uncertainty on the inputs, or specification of a problem, will affect the outputs or results. This is the problem of *propagation*. Part of this problem is to find which parameters contribute most to output uncertainty; these parameters will require the most effort to estimate in the *a priori* specification of the problem.

The main concern of this paper is the problem of uncertainty propagation and in particular, the propagation through nonlinear systems. Nonlinear systems present a specific problem in uncertainty analysis because they are not guaranteed to be structurally stable; they may bifurcate when small changes are made to the inputs or parameters. In contrast, linear systems are structurally stable and small perturbations of the inputs or parameters will induce small perturbations of the outputs. In fact this property provides the main line of attack for uncertainty propagation through linear systems; to quote [2], 'If the coefficients of variation of the involved quantities are small, then the perturbation method can be directly applied to the majority of problems...' (It is important to point out here that [2] is not subject to this restriction and proposes approaches which can deal with the situation when the coefficients of variation are moderate to large.)

It may actually be the case that the solution to one of the problems suggests or fixes the solution to one of the others; as an example consider the specification of uncertainty in terms of an *interval*. This is a situation in which nothing is known of a variable x except an upper bound \bar{x} and a lower bound \underline{x} ; the variable can then be expressed by the interval $[\underline{x}, \bar{x}]$. The adoption of this model for uncertainty, in a sense, fixes the strategy for uncertainty propagation. The reason is that there is a type of arithmetic for such uncertain quantities which mirrors all the usual operations for *certain* or *crisp* numbers; for example, addition of two interval numbers is governed by,

$$[a, b] + [c, d] = [(a + c), (b + d)] \quad (1)$$

Multiplication of two interval quantities is given by,

$$[a, b] * [c, d] = [\min(ac, ad, bc, bd), \max(ac, ad, bc, bd)] \quad (2)$$

(The latter definition takes into account the effect of negative values). Multiplication of an interval number by a positive scalar c is given by,

$$c * [a, b] = [ca, cb] \quad (3)$$

Finally, if an interval number is the input to a monotonically-increasing function F , then the result is given by,

$$F([a, b]) = [F(a), F(b)] \quad (4)$$

One can immediately see how uncertainty based on an interval prescription could be propagated through an engineering model, say a Finite Element (FE) model. A general crisp FE model is specified in terms of some initial quantities e.g. the mass and stiffness matrices for a dynamic calculation (the assembly of such matrices is in itself non-trivial, but that is another matter). A calculation can then proceed towards some desired output e.g. a set of natural frequencies. The algorithm for the calculation can ultimately be broken down into a sequence of the usual mathematical operations. Now, if the initial specification for the mass and stiffness matrices was in terms of intervals (matrices of intervals), in order to encode uncertainty as to their values, in principle, one could carry the calculation of interest forward by replacing each crisp operation by its interval analogue. The result of this procedure would (if successful) be an interval containing upper and lower bounds on the natural frequencies consistent with the uncertainties in the initial parameters or inputs to the model. In fact this type of approach, based on interval analysis, has been implemented with some success in terms of modal analysis and FE analysis [3]. One of the major problems with interval analysis is that it is inherently conservative; during long sequences of calculations, particularly if these are recursive, the uncertainty encoded in the width of the interval, can grow considerably. The final intervals from a long calculation are guaranteed to include the true range in the outputs engendered by the uncertainty in the inputs, but they are often much larger - so much larger in some cases, that they are useless. Another problem with interval analysis is that it represents the uncertainty in possibly the coarsest fashion imaginable. A much better theory of uncertainty is probability theory; this accounts for randomness in the most informative manner possible in direct contrast to interval analysis [4].

The specification of a (continuous) random variable X in probability theory is in terms of the probability density function (pdf) for the variable, $p(x)$. The usual definition of this quantity is that the probability that a random variable X will fall in the range $[x, x + dx]$ is given by,

$$P(X \in [x, x + dx]) = \int_{-\infty}^{\infty} p(x) dx \quad (5)$$

with the standard meanings for the various symbols. The pdf encodes all possible knowledge of the behaviour of X . So if probability is perfect, it begs the question - why not always use probability theory? The answer is, one does if one can, but there are practical issues which raise difficulties. The most immediate problem is that it is not always straightforward to obtain the pdf for a variable. If the pdf is not derivable from the underlying physics of the problem, it will need to be estimated or elicited from experts, and this can be a problem [4].

Suppose that the problem of interest is to compute an output y for a problem, given d input parameters $\{x_1, \dots, x_d\}$. For example, one may wish to know the maximum value of stress over a component modelled by FE with the various inputs being the material properties etc. which require prior specification. Regardless of the complexity of the calculation, it can be regarded as a function of the inputs,

$$y = f(x_1, \dots, x_d) \quad (6)$$

which results in the output of interest. So far, all of this is deterministic; however, suppose that the set of inputs is uncertain. If the inputs are random, the best specification one can have is the multivariate joint pdf (jpdf) $p(x_1, \dots, x_d)$, whose definition is a natural generalisation of (5). If the jpdf of the inputs is known, it may still be an intractable problem to translate this into a pdf (specification of uncertainty) for the output y . To explain things in the simplest terms, suppose that one has a univariate problem $y = f(x)$ (maximum stress only depends on Young's modulus, say). If the pdf of the input $p_x(x)$ is known, the pdf of the output $p_y(y)$ follows from a transformation law,

$$p_y(y) = p_x(f^{-1}(y)) \left| \frac{df^{-1}(y)}{dy} \right| \quad (7)$$

and this is all very well as long as one can write down the function f explicitly. If f is only specified implicitly by a large (potentially nonlinear) computer model, this will not be possible. A further complication is that very few problems indeed have a single input. In the absence of an analytical solution, one approach to the propagation problem is that of *Monte Carlo* (MC) analysis [5]. The idea is simply stated: one generates many random samples of possible inputs consistent with their jdpf and then propagates the individual crisp input sets through the algorithm or computation; the result is a set of outputs which one can use in order to estimate their pdf. The problem is that, if the input set is multi-dimensional, it may require very many samples to be drawn in order to cover all possibilities. This is fine as long as the calculation is quick; however, the implication is that one is dealing with a large complex calculation which may well be expensive to run in terms of computer resources. In the latter case, MC analysis may be prohibitively expensive.

One way out of this dilemma is to use a *surrogate model* or *emulator* (also called *metamodels*, *response surfaces*, *fast-running models* etc. under various circumstances). Again, the idea is simple to state: suppose the true (expensive) model specification is as in (6), a surrogate model is simply an alternative model $y_s(x_1, \dots, x_d)$ which behaves like the true system on a given set of inputs, but is structurally much simpler. The most basic examples of such a surrogate model might be a straight-line regression fit to a polynomial, or a polynomial fit to a transcendental function. The key idea here is exposed by the use of the term *regression*; the basis of regression analysis is that one can specify a parametric model for an input-output process and then *estimate* or *learn* the parameters of the model by minimising the model error on some known input-output data (usually called training data). This general regression problem is central to the well-explored discipline of *machine learning* [6]. The key observation is that the regression model can typically be learned using a small number of input-output sets, and this will require a smaller number of runs of the true model, certainly many less than a MC analysis would require. Once the surrogate model or emulator is established, it can then be used to generate the large number of input-output sets required by MC analysis in order to characterise the output uncertainty. A further critical advantage of the emulator is that, in contrast to the true model, it will often have a simple closed-form structure which could in principle be used for the analogue of equation (7), and this offers the possibility of avoiding MC completely. The idea of the emulator will be central to the next section's discussion of Bayesian sensitivity analysis.

Bayesian Sensitivity Analysis

The discussion of the last section was centred on the idea of uncertainty propagation; given a series of uncertain inputs to a computation, how does one characterise the uncertainty on the output? As discussed in the introduction there are other pressing problems in dynamics which are associated with uncertainty, one of them being *sensitivity analysis* (SA) [7]. In a sense, this is a refinement of the propagation problem where one wishes to determine (qualitatively) which input variables (with their associated uncertainties) are responsible for the output uncertainty. The SA technique which will be discussed and applied here is based on Bayesian probabilistic approach as detailed in [8]. (The approach is well-established within the probability and statistics community, the current paper is concerned with arguing the benefits of adopting it in the engineering context.) Each uncertain input parameter is represented as a probability distribution, and an emulator is fitted using multiple runs of the true model. In this case, the emulator is based on a powerful regression paradigm from machine learning called a *Gaussian process* [9]. From this emulator, statistical quantities relating to sensitivity and uncertainty can be inferred directly – for example, output uncertainty distributions and parameter sensitivities. A further advantage of using Gaussian process regression is that the uncertainty of the emulator fit is itself quantified.

Gaussian Process Regression. As discussed above, any computer model can be regarded as a function of its inputs: $f(\mathbf{x})$. (From this point, the multivariate input set will be specified as a

boldface vector; this releases subscript indices to indicate instances of training data.) Although the function is (usually) deterministic and governed by known mathematical relationships, it is often complex to the point of mathematical intractability. From a practical point of view, $f(\mathbf{x})$ can then be regarded as an unknown function, given that one does not know the output for a given set of inputs until one has actually run the model - or an emulator if the true model is computationally expensive. As indicated above, the emulator discussed here will be the *Gaussian process* (GP).

GPs are an extension of the standard multivariate Gaussian probability distribution, they are essentially a probability density function *over functions*. Whereas most forms of regression return a crisp value $f(\mathbf{x})$ for any given \mathbf{x} , a GP returns a specification for a Gaussian probability distribution. This means that prediction using the model automatically provides estimates of the confidence interval on the output. GPs also adhere to the Bayesian paradigm for probability, that is, a number of prior assumptions are made about the function being modelled, and then training data (samples from the model) are used to update and evaluate a posterior distribution over functions. A key assumption is that the model is a smooth function of its inputs – it is this that allows extra information concerning the response to be gained at reduced computational cost. Clearly if a function is smooth (in a sense, slowly-varying), information from a data point (input) will allow inferences on the behaviour of the function on neighbouring points.

The basic elements of GP regression will be discussed here, a much more extensive treatment can be found in [10]. For any set of n input point $\{\mathbf{x}_1, \dots, \mathbf{x}_d\}$ (the subscripts now provide a training point label), each of dimension d , the prior beliefs about the corresponding outputs (if a vector of outputs is considered, see later) is specified by a multivariate normal distribution. This in itself is not informative enough, complete specification of the Gaussian distribution requires some prior assumptions about its mean and covariance. The mean is assumed to have the form of a regression fit through the training data,

$$E\{f(\mathbf{x})|\boldsymbol{\beta}\} = \mathbf{h}(\mathbf{x})^T \boldsymbol{\beta} \quad (8)$$

where $\mathbf{h}(\mathbf{x})$ is a specified regression function of \mathbf{x} , and $\boldsymbol{\beta}$ is the corresponding vector of coefficients. The vertical line in Eq.8 indicates that this expression is a conditional expectation. For simplicity, $\mathbf{h}(\mathbf{x})$ is chosen here to be $(1, \mathbf{x})$, representing a linear regression,

$$E\{f(\mathbf{x})|\boldsymbol{\beta}\} = \beta_0 + \beta_1 x_1 + \dots + \beta_d x_d \quad (9)$$

(this can be extended to higher-order polynomial fits if required). Note that a straightforward regression approach would propose a relationship of the form,

$$y = f(\mathbf{x}) = \beta_0 + \beta_1 x_1 + \dots + \beta_d x_d \quad (10)$$

The coefficients for a standard least-squares regression model of the form (10) would usually be obtained by minimising the sum of the squared prediction errors over the training data; the result would be a rather minimal linear model in the context of the $\mathbf{h}(\mathbf{x})$ discussed here. Eq. (9) is rather more sophisticated than this for reasons which will become clear later. In addition to a prior specification for the mean for the GP model, one requires covariance information. The covariance between output points for the GP model is assumed to be,

$$\text{cov}\{f(\mathbf{x}), f(\mathbf{x}')|\sigma^2, B\} = \sigma^2 c(\mathbf{x}, \mathbf{x}') \quad (11)$$

where σ^2 is a scaling factor and B will be a diagonal matrix of length-scales, representing the roughness of the output with respect to the individual input parameters. These quantities are specified before any training on the data and are referred to as *hyperparameters*. The covariance function used here is chosen to be a squared-exponential function of the form,

$$c(\mathbf{x}, \mathbf{x}') = \exp\{-(\mathbf{x} - \mathbf{x}')^T B (\mathbf{x} - \mathbf{x}')\} \quad (12)$$

The posterior distribution for the GP is found by conditioning the prior distribution on the training data \mathbf{y} , and integrating out (or marginalising over) the hyperparameters σ^2 and $\boldsymbol{\beta}$. (Note that the output here is considered to be univariate and the vector is an ordered array of all of the output values in the training data; it is thus of dimensions $n \times 1$.) The process of marginalisation is central to the Bayesian approach; essentially one integrates the model over all possible values of the hyperparameters according to their probability. This means that the hyperparameters themselves require a prior specification for their probability density; (this will not be discussed further here, the curious reader can consult [12] for more detail). The integrals involved are usually all of a Gaussian form, and although the expressions are almost always very complicated, the results can be given in closed-form. The result is a Student's t-process, conditional on B and the training data,

$$[f(\mathbf{x}) | \mathbf{B}, \mathbf{y}] \sim t_{n-q}\{m^*(\mathbf{x}), \hat{\sigma}^2 c^*(\mathbf{x}, \mathbf{x}')\} \quad (13)$$

where,

$$m^*(\mathbf{x}) = \mathbf{h}(\mathbf{x})^T \hat{\boldsymbol{\beta}} + \mathbf{t}(\mathbf{x})^T A^{-1} (\mathbf{y} - H \hat{\boldsymbol{\beta}}) \quad (14)$$

$$c^*(\mathbf{x}, \mathbf{x}') = c(\mathbf{x}, \mathbf{x}') - \mathbf{t}(\mathbf{x})^T A^{-1} \mathbf{t}(\mathbf{x}') + (\mathbf{h}(\mathbf{x})^T - \mathbf{t}(\mathbf{x})^T A^{-1} H) (H^T A^{-1} H)^{-1} (\mathbf{h}(\mathbf{x}')^T - \mathbf{t}(\mathbf{x}')^T A^{-1} H) \quad (15)$$

$$\mathbf{t}(\mathbf{x})^T = (c(\mathbf{x}, \mathbf{x}_1), \dots, c(\mathbf{x}, \mathbf{x}_n)) \quad (16)$$

$$H^T = (\mathbf{h}(\mathbf{x}_1), \dots, \mathbf{h}(\mathbf{x}_n)) \quad (17)$$

$$A = \begin{pmatrix} 1 & c(\mathbf{x}_1, \mathbf{x}_2) & \dots & c(\mathbf{x}_1, \mathbf{x}_n) \\ c(\mathbf{x}_2, \mathbf{x}_1) & 1 & & \vdots \\ \vdots & & \ddots & \\ c(\mathbf{x}_n, \mathbf{x}_1) & \dots & & 1 \end{pmatrix} \quad (18)$$

$$\hat{\boldsymbol{\beta}} = (H^T A^{-1} H)^{-1} H^T A^{-1} \mathbf{y} \quad (19)$$

$$\hat{\sigma}^2 = \frac{1}{n-d-3} \mathbf{y}^T \{A^{-1} - A^{-1} H (H^T A^{-1} H)^{-1} H^T A^{-1}\} \mathbf{y} \quad (20)$$

$$\mathbf{y} = (f(\mathbf{x}_1), \dots, f(\mathbf{x}_n))^T \quad (21)$$

Note that the determination of the emulator (12) is basically an exercise in machine learning and as such its quality is critically dependent on the number and distribution of training data points in the input space, and the values of the hyperparameters. The diagonal matrix of roughness parameters B cannot generally be integrated out analytically and must be evaluated using *maximum a posteriori* (MAP) estimation – this calculation typically represents the most computationally intensive part of the process. It is useful to dwell a little on what has happened through the last sequence of equations. Arguably the most important part of the GP model is the mean of the model as this is what one takes as the prediction on a given \mathbf{x} . The prior specification of the model is given by Eq.(8) and carries no information yet about the training data. After incorporation of the training data (and integration over the hyperparameters), the posterior form of the mean is given by Eq.(14) and this can be written in a form which exposes its meaning a little more clearly; one writes,

$$m^*(\mathbf{x}) = \mathbf{h}(\mathbf{x})^T \hat{\boldsymbol{\beta}} + \Sigma(\mathbf{x}) \hat{\boldsymbol{\epsilon}} \quad (22)$$

where $\Sigma = \mathbf{t}(\mathbf{x})^T A^{-1}$ and $\hat{\boldsymbol{\epsilon}} = \mathbf{y} - \mathbf{H}\hat{\boldsymbol{\beta}}$. The latter term is essentially the vector of residual errors over the training set for a linear regression model. If the input-output process were truly linear and there was no measurement (or other) uncertainty, this term would be zero and Eq.(22) would collapse onto a perfect linear regression model. If the process is not linear and/or there is measurement uncertainty, the second term in Eq.(22) ‘switches on’ and modifies the linear regression core. The second term then acts as a sort of (non-polynomial) interpolant; it carries information about the training data through both terms and it also carries information about the assumed smoothness of the model through the matrix Σ which is built from data dependent on the covariance function. This combination of regression and interpolation in the posterior mean, means that it has much greater predictive power than a linear regression alone. In fact, the GP is even more sophisticated than this; as one also has a posterior variance for the model, one can establish natural confidence bounds on predictions.

The dependence of the emulator on training data means that some model runs are always required. As discussed above, the advantage of the Bayesian sensitivity approach is that, typically far fewer runs are needed to train the emulator than would be needed for a full MC analysis. The main reason for this is that the emulator carries information about the smoothness of the data and therefore knowing the function at one point in the input space means one has information about the likely values at neighbouring points; MC analysis does not allow for inference of this nature, the points at which the function is evaluated are considered isolated from each other. To deal with the sampling of training data in as principled a manner as possible, ideas of experimental design are applied and a *maximin Latin hypercube* design (maximin LHD) is used here [11]. A Latin hypercube design divides input space into regions of equal probability and distributes points in a *Latin-square* across ‘probability-space’. A maximin LHD improves on this by additionally maximising the minimum distance between input points, thus optimising the space-filling properties of the design.

Inference for Sensitivity Analysis. Having established an emulator, the next parts of the process are associated with establishing how uncertainty propagates through the emulator and finding the sensitivities of interest [8,9,12]. Several quantities can be inferred from the posterior distribution-over-functions described above, that are relevant to sensitivity analysis. Fundamental quantities such as the mean and variance of the output distribution can be evaluated and this provides an estimate of the output uncertainty. In addition to total output variance, one can determine *main effects*, *interactions* and other sensitivity measures for the input parameters, based on their contribution to, or responsibility for, output variance. If a small amount of uncertainty (low variance) on an input parameter generates a high degree of uncertainty (high variance) on the output, then the output is clearly sensitive to that parameter in terms of uncertainty propagation.

The function $f(\mathbf{x})$ can be decomposed as follows, into *main effects* and *interactions*,

$$y = f(\mathbf{x}) = E(Y) + \sum_{i=1}^d z_i(x_i) + \sum_{i<j} z_{ij}(x_i, x_j) + \dots + z_{12\dots d}(\mathbf{x}) \quad (23)$$

where the expectation operator E is defined here with respect to the probability distribution of the inputs and,

$$z_i(x_i) = E(Y|x_i) - E(Y) \quad (24)$$

$$z_{ij}(x_i, x_j) = E(Y|x_i, x_j) - z_i(x_i) - z_j(x_j) - E(Y) \quad (25)$$

with higher-order terms being defining in an analogous fashion. Here $z_i(x_i)$ represents the *main effect* of the input variable x_i ; $z_{ij}(x_i, x_j)$ is the *first order interaction* of the input variables x_i and x_j and further terms represent higher-order interactions; Y is the random variable corresponding to the function output y and therefore $E(Y)$ is the expected value of the output y considering all

possible combinations of inputs. The main effect of an input can be thought of as the effect (on the output) of varying that parameter over its input range, averaged over all the other inputs. Interactions describe the effect of varying two or more parameters simultaneously, additional to the main effects of both variables. Plotting main effects serves as a visual indication of the influence of particular inputs and interactions, showing (albeit qualitatively) the variance of the output with respect to individual input parameters and the nonlinearities associated with those responses. This will become clear when the case study material is discussed later.

One can infer posterior mean values for main effects and interactions by simply substituting the posterior mean into the definitions detailed above. The necessary tool is the conditional expectation defined by,

$$E(Y|\mathbf{x}_p) = \int_{\chi_{-p}} f(\mathbf{x})p(\mathbf{x}_{-p}|\mathbf{x}_p)d\mathbf{x}_{-p} \quad (26)$$

where the subscripts p and $-p$ indicate subsets of the inputs indexed by a set of labels p and the complement of p respectively and $p(\mathbf{x}_{-p}|\mathbf{x}_p)$ represents the multivariate probability density function of the input parameters indexed by $-p$ conditional on those indexed by p . The set χ_{-p} is that spanned by the input variables indexed by $-p$. The posterior mean of the main effect can be derived by substituting (26) into (14) in place of $f(\mathbf{x})$, the result is,

$$E^*\{E(Y|\mathbf{x}_p)\} = R_p(\mathbf{x}_p)\widehat{\boldsymbol{\beta}} + T_p(\mathbf{x}_p)\mathbf{e} \quad (27)$$

where E^* is the expectation with respect to the posterior distribution now (i.e. that governing the random variable Y , and,

$$R_p(\mathbf{x}_p) = \int_{\chi_{-p}} \mathbf{h}(\mathbf{x})^T p(\mathbf{x}_{-p}|\mathbf{x}_p)d\mathbf{x}_{-p} \quad (28)$$

$$T_p(\mathbf{x}_p) = \int_{\chi_{-p}} \mathbf{t}(\mathbf{x})^T p(\mathbf{x}_{-p}|\mathbf{x}_p)d\mathbf{x}_{-p} \quad (29)$$

$$\mathbf{e} = A^{-1}(\mathbf{y} - H\widehat{\boldsymbol{\beta}}) \quad (30)$$

Although this results in a series of matrix integrals, a Gaussian or uniform distribution over the input parameters allows these to be evaluated analytically. Expressions for interactions can be similarly derived using their respective definitions.

An alternative approach to summarising sensitivity information is in terms of *variance and sensitivity indices*; such methods are widely used [12]. This approach involves quantifying the proportion of output variance for which individual input parameters are responsible. In particular, sensitivity can be measured by conditional variance,

$$V_i = \text{var}\{E(Y|x_i)\} \quad (31)$$

This is the expected value of the contribution of the input variable x_i to the output variance. Note that this is also the variance of the main effect of x_i , hence it is known as the *main effect index* (MEI). This can be extended to measure conditional variance of interactions of inputs, i.e.,

$$V_{ij} = \text{var}\{z_{ij}(x_i, x_j)\} \quad (32)$$

and so on for higher-order interactions. Although this approach provides detailed insight into the effects of combinations of inputs on output uncertainty, it can be time-consuming to examine all possible interaction permutations for models with many input dimensions. An alternative sensitivity

measure [12], describes the output variance that would remain if one were to learn the true values of all inputs except x_i ,

$$V_{ti} = \text{var}(Y) - \text{var}\{E(Y|\mathbf{x}_{-i})\} \quad (33)$$

This last measure, called the *total sensitivity index* (TSI), measures the variance caused by an input x_i and any interaction of any order including x_i . It allows a more holistic view of the uncertainty attributed to each input, but does not give any details as to how it is distributed between main effects and interactions. Between the MEIs and TSIs, a detailed view of the sensitivities of inputs and their interactions can be gained.

To calculate posterior means and variances of the above quantities from the emulator, the following equation is obtained [12],

$$\begin{aligned} E^*(V_p) &= E^*(E\{E(Y|\mathbf{x}_p)^2\}) - E^*\{E(Y)^2\} \\ &= \hat{\sigma}^2[U_p - \text{tr}(A^{-1}P_p) + \text{tr}\{W(Q_p - S_p A^{-1}H - H^T A^{-1}S_p^T + H^T A^{-1}P_p A^{-1}H)\}] \\ &\quad + \text{tr}(\mathbf{e}^T P_p \mathbf{e}) + 2\text{tr}(\hat{\boldsymbol{\beta}} S_p \mathbf{e}) + \text{tr}(\hat{\boldsymbol{\beta}}^T Q_p \hat{\boldsymbol{\beta}}) - \text{var}^*\{E(Y)\} - \{E^*(E(Y))\}^2 \end{aligned} \quad (34)$$

where,

$$U_p = \int_{\chi_p} \int_{\chi_{-p}} \int_{\chi_{-p}} c(\mathbf{x}, \mathbf{x}^*) p(\mathbf{x}_{-p}|\mathbf{x}_p) p(\mathbf{x}'_{-p}|\mathbf{x}_p) p(\mathbf{x}_p) d\mathbf{x}_{-p} d\mathbf{x}'_{-p} d\mathbf{x}_p \quad (35)$$

and the terms P_p , Q_p and S_p have a similar integral structure. The symbol \mathbf{x}^* denotes a vector comprised of \mathbf{x}_p and \mathbf{x}'_{-p} . The presentation of these quantities is not just gratuitous, the motivation is to give some idea of the complexity of the sort of calculation that can nonetheless be carried out analytically in some cases of importance; this is one of the remarkable strengths of the Bayesian approach. The full details of the calculation of V_i and V_{Ti} can be found in [12]. All the quantities of interest presented here can be computed using the software package GEM-SA [11], and this was used in the case studies which now follow.

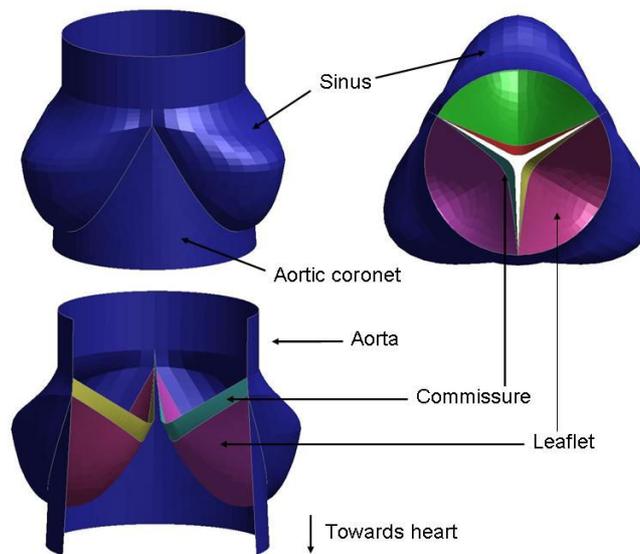


Figure 1. Layout of the aortic valve.

Case Study 1: A Heart Valve Model

Background. The aortic valve is situated in between the left ventricle and the aorta in the heart. Its function is to ensure one-way flow of oxygenated blood from the heart to the rest of the body. The valve consists of three leaflets and three sinuses, arranged as shown in Fig.1. When ventricular blood pressure is greater than aortic blood pressure (in systole), the leaflets open; conversely when the opposite is true (diastole) the leaflets are forced against each other, restricting any blood flow. In a full life span, the valve opens and closes around 3.7 billion times [13]. However, the aortic valve is prone to a variety of diseases and it is not uncommon that it requires replacement. Implanting a replacement that performs consistently under such demanding circumstances is a very challenging affair.

Modelling. A reasonable starting point in prosthesis design is to acquire a comprehensive understanding of the natural component to be replaced; one approach to this is via physical modelling. However, the uncertainty involved in modelling a heart valve is substantial – geometry, material properties and loading can only be specified to within fairly wide ranges. Furthermore, the complicated nature of the material properties is an additional source of epistemic uncertainty - indeed, some work has been done to investigate this [14]. Clearly a comprehensive uncertainty analysis of a typical heart valve FE model would provide useful insight into the robustness of the simulation. For the purposes of this study, a 3-dimensional dynamic FE model of the heart valve was created parametrically in ANSYS and solved in the LS-Dyna package. The approach was to use Belytschko-Lin-Tsai shell elements to model the entire structure, encompassing the sinus, leaflets and a section of aorta at either end. An illustration of the mesh is shown in Fig.2. The simulation was performed over the initial opening movement of the valve: a total time of 0.1s, starting from the point where the pressure on the leaflet becomes positive in the z-direction. The model consisted of around 3200 elements, depending on the set of input parameters used.

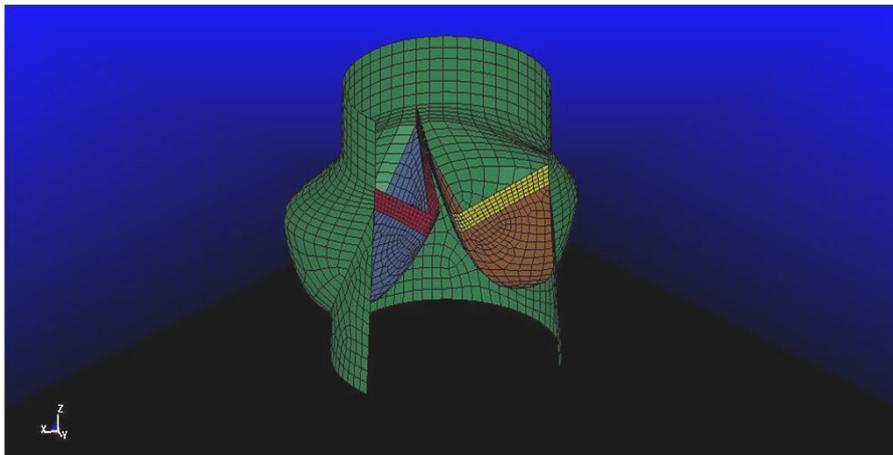


Figure 2. Two-thirds of aortic valve mesh (typical parameter set).

The geometry is described with a relatively small number of parameters [13], each of which is quoted as a range. Some finer geometrical points in the model are left to the modeller's discretion, as describing exact dimensions of a naturally varying biological system is meaningless. The thickness of the tissue in the leaflet also varies from commissures to leaflet tips, however it was considered sufficient in this investigation to have only two regions of different thickness – one for the commissures and the other for the rest of the leaflet.

The material of the heart valve is anisotropic and nonlinear, and is different between the leaflet and the sinus. The leaflet is comprised of a series of collagen fibres running largely in the circumferential direction, joined by loose connective tissue. At low strains these fibres are not under tension and resistive force is due purely to the connective tissue, however for greater extensions the

fibres provide a stronger elastic resistance – this is responsible for the hyperelastic nature of the material. The aortic tissue is also hyperelastic, but due to the largest strains occurring in the leaflet, it is usually considered sufficient to model it as linear. Although both leaflet and aortic tissue is hyperelastic and anisotropic, for the purposes of this investigation the material was assumed to be elastic, albeit having separate elastic moduli between leaflet, commissure and sinus. One justification for this is that existing uncertainty data on distributions of elastic moduli used for modelling aortic valves is scarce enough – hyperelastic parameters even more so. Furthermore, hyperelastic material models introduced too many input dimensions for the scope of the study.

A number of recent aortic valve simulations have made use of fluid-structure interaction techniques to more realistically model blood/leaflet interaction [15,16], however in the interests of reducing computational expense, for this investigation the loading was applied as a set of simple pressure curves, assumed to be uniform across the leaflet [16]. The respective pressures on the leaflet, and the aortic and ventricular sides of the valve were obtained from the pressure curves in the cardiac cycle, which are well established quantities. Although this is not a completely accurate representation of the leaflet loading (since this is an interactive relationship that depends on the position of the leaflets), it was considered adequate for the purposes of this uncertainty analysis.

Table 1. Chosen Input parameters for sensitivity analysis.

Parameter	Name	Mean	Standard Deviation	Data source
Elastic modulus of sinus (MPa)	E_s	2.45	0.8	Lit. Review
Thickness of sinus shells (mm)	T_s	1.3	0.325	Lit. Review
Elastic modulus of leaflet (MPa)	E_l	2.05	0.65	Lit. Review
Thickness of leaflet shells (mm)	T_l	0.35	0.0085	Lit. Review
Thickness of commissure shells (mm)	T_c	1	0.25	Lit. review/Thubrikar
Radius of base (mm)	R_b	12.7	0.675	Thubrikar
Angle of coaptation (degrees)	Φ	31	3	Thubrikar
Leaflet separation in relaxed state (ratio)	$Fltrad$	6	2	Subjective

Sensitivity Analysis. A full holistic sensitivity analysis of a FE model is impractical, even with the Bayesian approach outlined above, since the number of uncertain parameters (geometrical, material, loading and numerical) can be extended almost without limit. It was therefore considered enough to pick eight input parameters to examine, that were found to vary substantially from one simulation to the next in a comprehensive literature review on the topic. Ranges could then be collected for each parameter based, in a loose sense, on the “belief” expressed in the existing literature, in the absence of formal statistical data. Other parameters were collected from studies quoted in [13]. The input parameters are presented in Table 1. Each parameter was assigned a normal distribution for mathematical tractability.

Only a brief illustration of the results of the analysis can be given here, the reader is referred to [17] for much more detail. The methods used in this investigation require that the output of the model be a scalar quantity; this meant that a certain subjectivity crept into the analysis, a prior selection of a small number of quantities of interest was made and a separate sensitivity analysis was carried out for each. The output parameters considered are shown in Table 2, representing: maximum stress, speed of valve opening (which is one of the methods used to validate aortic valve models [18]), maximum displacement of leaflets and leaflet buckling (which appears to important in the mechanisms of valve failure [19]). In each case, the input ranges were used to create a maximin Latin hypercube design of 250 training points. This input matrix was then run through LS-Dyna to

evaluate the model 250 times and provide corresponding output points. Separate sensitivity analyses were then run for each output of interest using GEM-SA, which generates main effect plots as well as MEIs and TSIs.

Table 2. Output parameters chosen for analysis.

Output Parameter	Description
<i>Sigmaxmax</i>	Maximum stress in any part of the model, at any time step
<i>Thalfop</i>	Time when leaflet tips pass between coaptation points (point at which leaflet is likely to buckle most)
<i>Topen</i>	Time for leaflet tips to reach aortic radius
<i>Dispmax</i>	Maximum displacement of leaflet tips (vector sum)
<i>Wig</i>	Percentage buckling at <i>Thalfop</i> (defined by angles successive elements make with each other in the x-y plane)

Results. The posterior expected value of the mean and variance for each output parameter is shown in Table 3. This gives an overall idea of the uncertainty in each output, given the uncertainty assigned to each input. To illustrate the comparative uncertainty between outputs, the standard deviation has also been presented as a percentage of the mean. It is evident that some of the uncertainties here are substantial – *Thalfop*, which represents the speed of opening of the valve, varies over 100% inside the 95% confidence limit. Similarly, the maximum stress value varies by around one order of magnitude.

Table 3. Output uncertainties.

Parameter	Posterior mean (μ)	Posterior standard deviation (σ)	σ/μ (%)
<i>Dispmax</i>	0.0125	0.00160562	12.8
<i>Thalfop</i>	0.00194	0.00122988	63.4
<i>Sigmaxmax</i>	635878	354868.99	55.8
<i>Wig</i>	0.161268	0.01573881	9.8

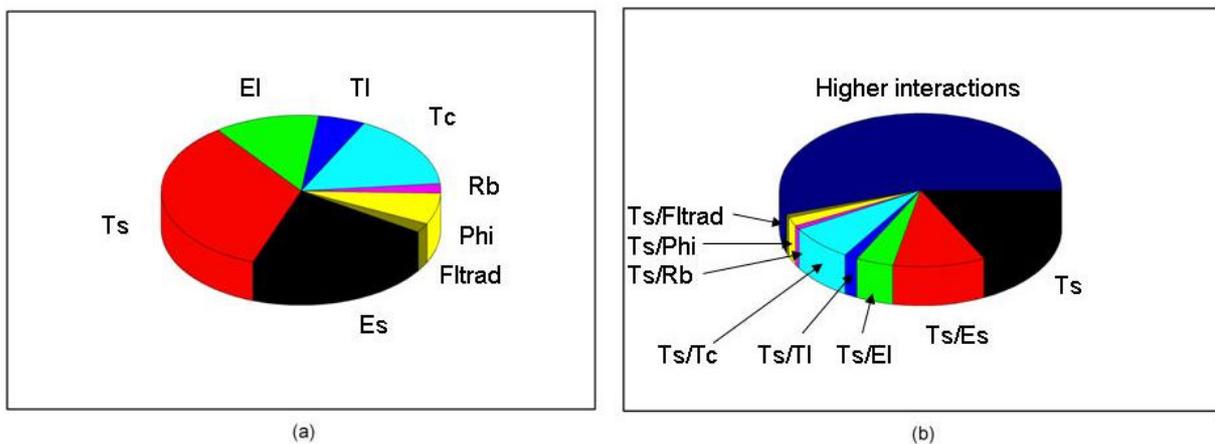


Figure 3. *Thalfop* sensitivity indices: (a) Total effect indices; (b) Interactions involving *Ts*.

A more detailed analysis shows that the uncertainties can be the result of high-order interactions. For example, less than one third of the uncertainty in the opening time (*Thalfop*) is explained by

first-order interactions. An examination of the total effect indices gives more insight, as illustrated in Fig.3a. Notably, the thickness and elastic modulus of the sinus and their interactions with other variables (Fig.3b) are responsible for a large proportion of the output variance – this supports the supposition that the flexibility of the sinus is partly responsible for the ease of opening of the valve [19]. Another first-order effect worth mentioning is that interactions between sinus thickness and commissure thickness play an important role (7.5%), possibly because the thicker commissure needs a flexible sinus to allow it to reverse its curvature.

In order to illustrate the information in the main effects, Fig.4 shows a main effect plot for the leaflet stiffness and reveals that there is a significant nonlinearity in the response of *wig* to *E1*. It is evident that below an elastic modulus of around 500 kPa the level of bucking increases dramatically.

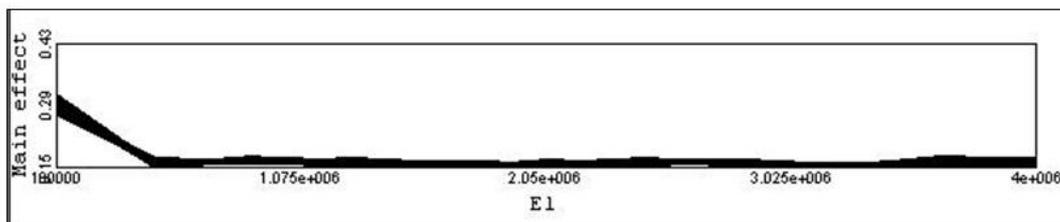


Figure 4. Main effect plot of *E1* for output *Wig*.

Discussion. Only a small fraction of the available results have been given here; even so, one can see the value of the analysis. Table 1 shows that for the four parameters investigated there is very substantial variation as a result of input uncertainty, with some standard deviations being over 50% of the mean. It should be remembered also that this is only due to the uncertainty in 8 input parameters, whereas in reality many more inputs are subject to variance, suggesting that these estimates are if anything, slightly conservative. This highlights a general problem in modelling biomechanical systems – often data is scarce, and aside from that, aleatoric uncertainties are unavoidable. At best, quantitative results from this model could only be quoted to quite vague ranges, given the assumed input uncertainty. The sensitivity analyses also show that there are likely to be complex interactions at work in a typical heart valve model.

Case Study II: Fluid Structure Interaction - An Unmanned Airship Model

For a range of aerial applications, lighter-than-air aircraft – *airships* – are enjoying a resurgence in popularity, in particular in the unmanned version. In order to combat some of the problems encountered by conventional airship designs and expand the range of operable weather conditions, a new design has been proposed in [20], called the Elettra Twin Flyers (ETF) (Fig.5). The design innovations are twofold – firstly there are two gas envelopes rather than one, positioned side by side. This allows a smaller profile for the equivalent lift. The two balloons are connected by a rigid structure that also acts as a platform for affixing a payload (e.g. surveillance equipment) and operational equipment. Secondly, the ship is moved in all six degrees of freedom by a series of directional propellers; the advantage being that the ship can be manoeuvred effectively without the requirement of forward motion. The innovations are intended to give the ship greatly increased manoeuvrability at low velocities, and allow the ETF to hover with the prow orientated in any direction, even in adverse weather conditions. This case study is concerned with the structural analysis of a model corresponding to the prototype design of the ETF.

The model is complex and only the briefest summary of its salient features will be given here, the reader is referred to [21] for much more detail (that reference also surveys the relevant background literature on the modelling of airships). In terms of overall scale, the model investigated here is 36m in length, this being the length required for sufficient buoyancy to support the total weight of the structure and payload (approximately 3 tonnes).

Material. In a sense, the material properties of the structure are not of great interest here as the main uncertainty discussed later will relate to the direction of aerodynamic loading; however, they are mentioned here as the materials used are quite novel and are themselves subject to uncertainty, and this is important for static stress analysis [21]. The frame material is a sandwich layup, consisting of layers of carbon/epoxy T300 15k/976 composite, either side of an Ultracor® honeycomb core [22]. This has the benefit of low density and high stiffness and resulted in a quasi-isotropic material at the macroscopic level, to simplify modelling. However, it was found in [23], that the material properties of the composite varied significantly with changes in temperature. Since the airship is potentially operating in conditions of $\pm 50^{\circ}\text{C}$ (dependent on altitude, location and weather), the uncertainties in material properties were expected to have a substantial effect on the response of the structure to loading.

Airship Model

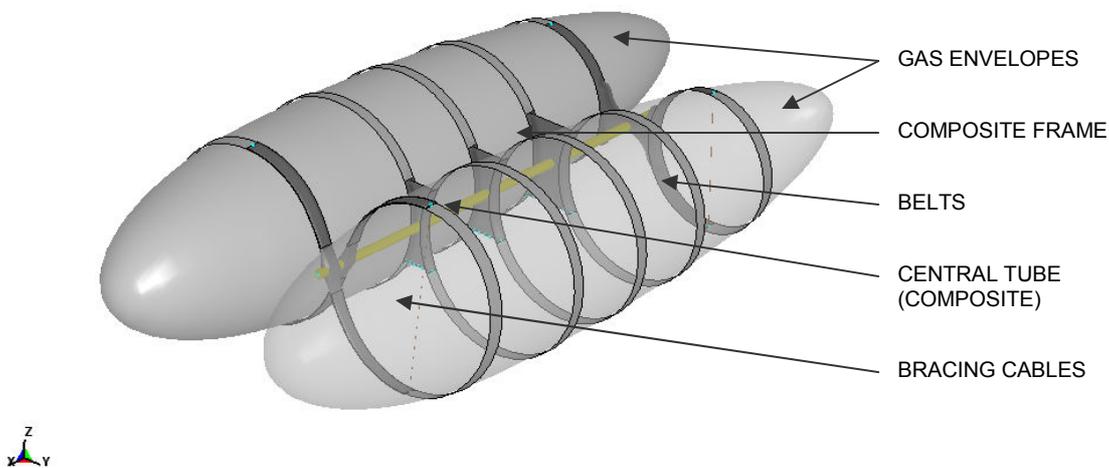


Figure 5. Illustration and arrangement of materials in the ETF

Modelling. The modelling of the airship was performed to reveal the stress and deformation of the airship under aerodynamic loading conditions. Control of the structure is currently based on the assumption that the structure is rigid; this is an important point because even minor deformations and rotations of the arms on which the driving propellers are located can have effects on the control system that are hard to predict. One main objective of the FE analysis was thus to quantify deformation in certain parts of the structure under a variety of loading conditions. In terms of structural integrity, both linear and nonlinear static models were constructed in the programme of study in order to study structural stresses, and the results are presented in [21]; however, the results shown here are from a full Fluid Structure Interaction (FSI) model considered the effect of gusts of wind incident at different angles to the airship. As in the case of the heart-valve model, the FE software LS-Dyna was used; the solver being capable of handling highly nonlinear and transient simulations, and the coupling of Lagrangian and Arbitrary-Lagrangian-Eulerian (ALE) meshes required for FSI.

Aerodynamic Loading. The forms and magnitudes of aerodynamic loading can vary greatly, from the response of the structure in steady-state flow, to ramped or stepped loading (i.e. from gusts of wind in still air). The gust is an important aspect in the study of the interaction of any aircraft with the real atmosphere and was considered to be crucial in this study. The form of gust model used is the classic constant-gradient gust referred to for example in the European Aviation Safety Agency regulations [24]. The velocity profile of this linear gust is known to closely follow that of a natural gust. Since the pressure distribution on a non-trivial structure due to a particular gust of wind is far

from obvious (especially at certain angles), the FSI model was created with an Arbitrary-Eulerian-Lagrangian (ALE) mesh overlapped with the Lagrangian mesh used for the “dry” model. The possibility of structural deformation that would change the boundary conditions of a standard CFD model necessitates this approach. The ALE mesh combines the benefits of the Lagrangian and Eulerian mesh approaches, respectively that the mesh is allowed to move in space (thus limiting the size of the mesh necessary); and that the material can flow through the mesh, allowing the very large deformations that are characteristic in fluid flow. An extensive discussion on the ALE method is found in [25]. The two meshes were coupled together using a penalty-based approach [26], Fig. 6 illustrates the arrangement of the two meshes. The ALE mesh is cylindrical because it allows investigation of different incident gust angles; the airship can simply be rotated in the mesh without the need for extra pre-processing.

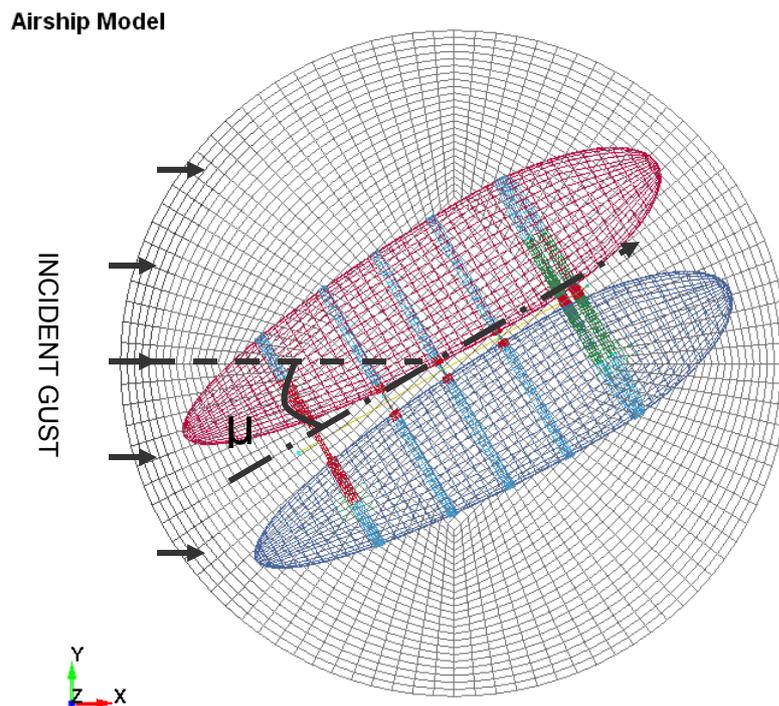


Figure 6. Illustration of cylindrical ALE mesh overlapping Lagrangian airship mesh, viewed from above; μ = angle between gust direction and ship axis.

The air is modelled in terms of compressible, viscous flow without considering turbulence; this represents a first approximation of the airflow around the structure for use in uncertainty analysis. Although there are many conceivable uncertainties associated with the definition of the gust hitting the airship, in this analysis only the wind angle was investigated. After adding the fluid mesh, the run time of the model became very substantial, so for this stage of the investigation the other uncertainties were treated as known to save on analysis time, though they may be investigated in future analyses.

Results and Discussion. To show the effects of the gust direction on an output of the model, Fig. 7 (left) shows the main effects on final displacement between opposite propellers in the x-direction of the model, (this is an important quantity, being a measure of the deformation of the structure and thus potential degradation of the control system). Extracting a meaningful quantity from this model is not trivial because the impact of the gust causes structural oscillations, and all displacement and stress values are a function of time. The variation in displacement is reasonably linear and quite small. This is attributed to the rigidity of the structure caused by tension in the balloons. The

displacement is increasingly negative as the gust is angled more towards the front of the airship, likely because the balloon is pushed towards the rear of the ship, causing the structure to contract slightly.

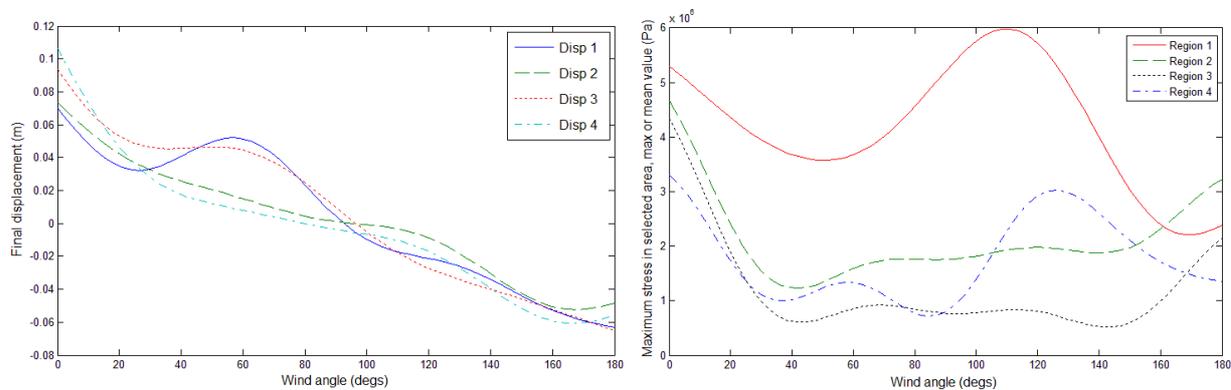


Figure 7. Posterior means of: displacements 1-4 (left), stresses in regions 1-4 (right).

Other outputs of the model are less well-behaved. A significant difficulty with this model is that a small change in gust angle can drastically change the output of the model. Fig.7 (right) shows the variation of stress with gust angle for four chosen regions in the model. Each of these regions represents an area of observed stress concentration in the frame. The response of stress is shown to be highly nonlinear in relation to the gust angle. This can sometimes cause problems with an emulator approach since a highly-structured response requires a higher density of training data (however, a check of the emulator accuracy using cross-validation showed that it was acceptable here). Finally, in a model like the one presented above, the number of uncertainties can be extended almost without limit. Loading uncertainties such as wind pressure distribution, velocity, air density and viscosity are could all vary (in fact, co-vary) due to environmental conditions or natural variation. Considering all these uncertainties in such a model is currently impractical, but it is worth remembering that considering only a small set of the model uncertainties will result in an underestimate of the true output uncertainty.

Summary

It does not seem appropriate to present a long summary here. The objective of the paper was simply to give an overview of some issues relating to uncertainty propagation through large (computationally-expensive) structural models. In fact, much of the discussion centred on the more refined problem of sensitivity analysis. The main thesis of the paper is that an established Bayesian approach to the sensitivity analysis, based on a Gaussian process emulator or surrogate model, can yield insight into the problem without the restrictive requirement of performing many model runs. (More correctly, the approach needs many less model runs than Monte Carlo.) In order to illustrate the sort of information which is readily available from the Bayesian analysis, two case studies were presented. Between the two case studies, one sees many of the aspects of engineering models which present difficulties for the modeller.

Acknowledgements

The authors would like to sincerely thank their collaborators on various aspects of the research presented here: Dr Jen Rowson, Dr Graeme Manson and Dr Jeremy Oakley of the University of Sheffield; Dr Alaster Yoxall of Sheffield Hallam University and Professor Piero Gili of the Politecnico di Torino. The authors would like to thank Dr Mark Kennedy for the use of the GEM-SA software and Ove Arup for the use of LS-Dyna. The activity regarding the airship model presented in this article is part of the project entitled ‘‘Innovative Solutions for Control Systems,

Electric Plant, Materials and Technologies for a Non-Conventional Remotely-Piloted Aircraft” funded by the Piedmont Region of Italy. The heart-valve model work is part of the research from the Sheffield Dynamics Group EPSRC Platform Grant and the funding from EPSRC is gratefully acknowledged.

References

- [1] G.J. Klir, and R.M. Smith: *Annals of Mathematics of Artificial Intelligence* Vol.32 (2001) pp.5-33.
- [2] I. Elishakoff, and Y. Ren: *Finite Element Methods for Structures with Large Stochastic Variations* (Oxford Texts in Applied and Engineering Mathematics, 2003).
- [3] B. Lallemand, A. Cherki, T. Tison and P. Level; *Journal of Sound and Vibration* Vol.220 pp.353-364.
- [4] A. O'Hagan and J.E. Oakley; *Reliability Engineering and System Safety* Vol.85 (2004) pp.239-248.
- [5] Y.A. Shreider; *Method of Statistical Testing: Monte Carlo Method* (Elsevier, Amsterdam; London, (1964).
- [6] C.M. Bishop; *Pattern Recognition and Machine Learning* (Springer-Verlag, 2007).
- [7] A. Saltelli, K. Chan and E.M. Scott; *Sensitivity Analysis* (Wiley, New York, 2000).
- [8] J.E. Oakley, and A. O'Hagan; *Journal of the Royal Statistical Society, Series B* Vol.66 (2004) pp.751-769.
- [9] M.C. Kennedy, C.W. Anderson, S. Conti and A. O'Hagan; *Reliability Engineering and System Safety* Vol.91 (2006) pp.1301-1309.
- [10] C.E. Rasmussen and C.K.I. Williams; *Gaussian Processes for Machine Learning* (MIT Press, 2006).
- [11] *Gem-SA Homepage*. <http://ctcd.group.shef.ac.uk/gem.html> (accessed: 27/03/10).
- [12] T. Homma and A.K. Saltelli; *Reliability Engineering and System Safety* Vol.52 (1996) pp.1-17.
- [13] M. Thubrikar; *The Aortic Valve* (CRC Press, Boca Raton, Fla., 1990).
- [14] A. Ranga, R. Mongrain, R. Mendes Galaz, Y. Biadillah and R. Cartier; *Journal of Medical Engineering and Technology* Vol.28 (2004) pp.95-103.
- [15] J. De Hart, G.W. Peters, P.J. Schreurs, and F.P. Baaijens; *Journal of Biomechanics* Vol.36 (2003) pp.103-112.
- [16] C.J. Carmody, G. Burriesei, I.C. Howard and E.A. Patterson; *Journal of Biomechanics* Vol.39 (2006) pp.158-169.
- [17] W. Becker, J. Rowson, J. Oakley, A. Yoxall, G. Manson and K. Worden; in *Proceedings of 23rd International Conference on Noise and Vibration Engineering, Leuven* (2008) on CD..
- [18] I. Vesely, D. Boughner and T. Song; *Annals of Thoracic Surgery*, Vol.46 (1988) pp.302-308.
- [19] R.J. Brewer, J.D. Deck, B. Capati and S.P. Nolan; *Journal of Thoracic and Cardiovascular Surgery*, Vol.72 (1976) pp.413-417.
- [20] M. Battipede and P.A. Gili; *Peculiar Performance of a New Lighter-Than-Air Platform for Monitoring*, 2004 Aviation Technology, Integration and Operation Forum, Chicago, IL – USA, September 2004.
- [21] W.E. Becker, K. Worden, M. Battipede and C. Surace; *Submitted to AIAA Journal* (2010).
- [22] A. Cappadona; *Analisi Strutturale su Dirigibile Non Convenzionale a Controllo Remoto: Configurazione Esoscheletro*, in *Facoltà di Ingegneria, Politecnico di Torino* (2008).
- [23] Defense, U.S.D.o., *Polymer Matrix Composites Material Properties*, in *Composite Materials Handbook*. (2002).
- [24] EASA, *Certification Specifications for Normal, Utility, Aerobatic and Commuter Category Aeroplanes CS23, Book 1*, E.A.S. Agency, Editor. (2009).
- [25] J. Donea, A. Huerta, J.-Ph. Ponthot and A. Rodriguez-Ferran; *Arbitrary Lagrangian-Eulerian Methods*, in *Encyclopedia of Computational Mechanics*, E. Stein, R. De Borts, T.J.R Hughes (Eds.) (John Wiley & Sons, New York, 2004).
- [26] J.O. Hallquist; *LS-Dyna Theory Manual* (Livermore Software Technology Corporation, 2006).