

promoting access to White Rose research papers



Universities of Leeds, Sheffield and York
<http://eprints.whiterose.ac.uk/>

This is an author produced version of a paper due to be published in **Mechanical Systems and Signal Processing**.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/43225>

Published paper

Fricker, T.E., Oakley, J.E., Sims, N.D., Worden, K. (2011) *Probabilistic uncertainty analysis of an FRF of a structure using a Gaussian process emulator*, Mechanical Systems and Signal Processing (In Press)
<http://dx.doi.org/10.1016/j.ymssp.2011.06.013>

Probabilistic uncertainty analysis of an FRF of a structure using a Gaussian process emulator

Thomas E. Fricker^a, Jeremy E. Oakley^b, Neil D. Sims^c, Keith Worden^d

^a*Corresponding author.*

*Department of Probability and Statistics, University of Sheffield,
Sheffield S3 7RH, UK*

Tel: +44 (0)114 222 3801. Fax: +44 (0)114 222 3809.

Email: t.fricker@shef.ac.uk.

^b*Department of Probability and Statistics, University of Sheffield,
Sheffield S3 7RH, UK.*

Email: j.oakley@shef.ac.uk

^c*Department of Mechanical Engineering, University of Sheffield,
Sheffield S1 3JD, UK.*

Email: n.sims@sheffield.ac.uk

^d*Department of Mechanical Engineering, University of Sheffield,
Sheffield S1 3JD, UK.*

Email: k.worden@sheffield.ac.uk

Abstract

This paper introduces methods for probabilistic uncertainty analysis of a frequency response function (FRF) of a structure obtained via a finite element (FE) model. The methods are applicable to computationally expensive FE models, making use of a Bayesian metamodel known as an *emulator*. The emulator produces fast predictions of the FE model output, but also accounts for the additional uncertainty induced by only having a limited number of model evaluations. Two approaches to the probabilistic uncertainty analysis of FRFs are developed. The first considers the uncertainty in the response at discrete frequencies, giving pointwise uncertainty intervals. The second considers the uncertainty in an entire FRF across a frequency range, giving an uncertainty envelope function. The methods are demonstrated and compared to alternative approaches in a practical case study.

Keywords: finite element model, probabilistic uncertainty analysis, envelope frequency response function, Gaussian process, metamodel, Bayesian

1. Introduction

Finite Element (FE) modelling is perhaps the most widely used computational tool in the analysis of structural vibrations, particularly for the prediction of frequency response functions (FRFs). In recent years there has been a growing level of interest in how different types of uncertainty can be handled with this modelling approach. These uncertainties can be inherent to the model itself

(for example due to assumptions regarding the boundary conditions), or alternatively they could arise due to unknown values of physical parameters (for example component geometry or material properties). In the latter case, this lack of knowledge could be attributed to variation between nominally identical components (i.e. variability), or uncertainty during the design process regarding the final choice of dimensions or material.

There has long been interest in how uncertainty propagates through FE models. The method with the greatest pedigree is the Stochastic Finite Element method (SFE) [1]; this is a probabilistic method. In the general SFE formulation, the material properties across the structure can be specified as a random field. In a manner similar to the discretisation of the structure into finite elements, the random field is discretised into a denumerable set of random variables using the Karhunen-Loeve expansion, which is then truncated at some finite order. The results from the FE model are then expressed as a mean value supplemented by an expansion in terms of the random variables, allowing statistics of the quantity of interest to be computed. In the last decade, interest has grown in possibilistic approaches, such as a fuzzy approach to FE analysis and computation of modal quantities [2, 3]. More recent work has considered component mode synthesis as a framework for investigating both probabilistic and possibilistic uncertainties [4], and using possibilistic techniques based upon fuzzy numbers [5]. A ‘fuzzy FE’ approach is also developed in [6] and applied to a variety of case studies including the Garteur benchmark FE problem - a small scale aircraft model developed for assessing ground vibration test techniques.

A common issue when trying to propagate uncertain parameters through complex FE models is that the deterministic nature of the modelling approach leads to many model evaluations being performed, each for a different configuration of the uncertain inputs. In probabilistic modelling, this results in Monte Carlo simulations, whilst in possibilistic modelling, the repeated model evaluations can be used to generate fuzzy numbers representing the uncertainty in the model’s response.

This paper focusses on a probabilistic method for uncertainty analysis of FE models, using a statistical metamodel, or emulator, to reduce the number of FE model evaluations required, and hence reduce the computational cost. The remainder of the paper is organised as follows. First, the probabilistic uncertainty analysis problem is formulated, before introducing the concept of the emulator. Next, the use of the emulator for uncertainty analysis is described using a simple graphical example. The uncertainty analysis of FRFs that are predicted from FE models is then considered. This approach is then applied to a numerical case study based upon the Garteur testbed. Following a discussion, conclusions are drawn regarding the application of this modelling approach to FE modelling problems in structural dynamics.

2. Probabilistic uncertainty analysis of FE models

Consider a deterministic FE model evaluated at a particular degree of freedom. It takes a set of p input parameters, denoted $\mathbf{x} = (x_1, \dots, x_p)^T$, and

returns a set of outputs that consists of pairs of modal parameters. A typical FE analysis considers a subset of the modal parameters, which we denote $\{(\hat{m}_i, \hat{k}_i) : i = 1, \dots, n_{modes}\}$, where \hat{m}_i are the modal masses and \hat{k}_i are the modal stiffnesses. The FE model is deterministic, so repeated runs with the same configuration of input parameters will return the same outputs, and we may represent it as a function $\mathbf{y} = \boldsymbol{\eta}(\mathbf{x})$, where $\mathbf{y} = (\hat{m}_1, \hat{k}_1, \hat{m}_2, \hat{k}_2, \dots, \hat{m}_{n_{modes}}, \hat{k}_{n_{modes}})^T$.

In the probabilistic uncertainty analysis of a deterministic computer model, we consider the values of the uncertain input parameters to be a multivariate random variable X . As a result, the output of the model is also a multivariate random variable, which we denote $Y = \boldsymbol{\eta}(X)$. The first step in the analysis is to quantify the uncertainty in X by specifying a probability distribution $F(\mathbf{x})$. This distribution may be constructed using data, or by eliciting expert opinion [7], or a combination of both. Our aim is then to propagate the uncertainty in X through the computer model in order to characterise the distribution of Y , which is known as the *uncertainty distribution*.

A straightforward solution to this problem is to use a Monte Carlo procedure. In this we draw a large sample $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ from the input distribution $F(\mathbf{x})$ and run the model at each sampled input configuration \mathbf{x}_i . The result is a sample of the outputs $\{\mathbf{y}_1, \dots, \mathbf{y}_N\}$, from which we can estimate any summary of the uncertainty distribution such as the mean, the variance, or a particular quantile, using the corresponding summary statistic. For example, the mean of the uncertainty distribution may be estimated using the sample mean of $\{\mathbf{y}_1, \dots, \mathbf{y}_N\}$. For a general summary, denoted $S(Y)$, the precision of the estimate is determined by the sample size N , and standard techniques are available for estimating the Monte Carlo error in the estimate [8].

When we perform an uncertainty analysis of an FE model of a structure, characterising the uncertainty distribution of the FE model outputs (i.e. the modal parameters) is often only an interim step. In many cases, we are ultimately interested in quantifying the uncertainty in the corresponding FRF of the modelled structure. According to the concept of modal superposition, the FRF of the undamped structure is calculated as

$$G(\omega; \mathbf{y}) = \sum_{i=1}^{n_{modes}} \frac{1}{\hat{k}_i - \omega^2 \hat{m}_i}. \quad (1)$$

Since the modal parameters are uncertain, the FRF at a particular frequency ω is itself a random variable, which we denote G_ω . Given the Monte Carlo sample of the modal parameters, we may obtain a sample from the distribution of the FRF at ω by simply plugging the sampled modal parameters into Equation (1). This gives us a sample $\{G(\omega; \mathbf{y}_1), \dots, G(\omega; \mathbf{y}_N)\}$ from which we may obtain any summary of the FRF uncertainty distribution, $S(G_\omega)$.

In the next section, an alternative approach is described based upon the use of an emulator. However, at this stage it is useful to briefly mention the practical relevance of Equation 1. Real structures possess damping, such that the modal solution involves an imaginary term. Nevertheless, most FE solutions do not

consider structural damping, and so any uncertainty in the damping does not directly influence the problem of uncertainty propagation in the FE analysis. Another aspect of Equation 1 is that there are more generalised modal solutions that involve mass-normalised modes and modal constants, rather than mass and stiffness terms. Furthermore, uncertainty can cause the density functions for the natural frequencies to overlap and give a finite probability that mode i will appear at a higher frequency than mode $i + 1$. This means that the emulator cannot distinguish between individual modes based upon their natural frequency. These issues will not be considered in the present study, since the intention here is to demonstrate that multivariate emulators can be applied to the uncertain FE problem in its simplest form, without introducing additional levels of complexity.

3. Emulators

Monte Carlo uncertainty analysis requires the model to be run at many input configurations in order to make accurate inference about the uncertainty distribution, and the number of runs required increases exponentially with the number of uncertain input parameters. Consequently, the Monte Carlo method described above is impracticable if the model is computationally expensive. A solution to this problem is to use a *metamodel*. A metamodel is a surrogate which mimics the behaviour of the model while being computationally cheap to run. It is trained using a small number of model runs, then used as a replacement for the model in the Monte Carlo procedure.

A metamodel may be statistical or non-statistical. Techniques used for building non-statistical metamodels include neural networks [9], support vector machines [10, 11], and genetic programming [12]. For a statistical metamodel, any statistical regression technique may be used. Popular approaches include response surface methodology, in which the model is represented by a low degree polynomial [13, 14], and nonparametric regression using Gaussian processes (GPs) [15]. GP regression also appears in geostatistics, where it is known as *kriging* [16, 17], and in the field of machine learning [18].

In this paper we use statistical metamodels, taking a Bayesian approach. We consider the model to be an unknown function $\boldsymbol{\eta}(\cdot)$ and give it full probabilistic specification in the form of a prior distribution. The prior distribution is updated using the training data, resulting in a metamodel (the posterior distribution) that gives an estimate of $\boldsymbol{\eta}(\cdot)$, but also quantifies uncertainty about $\boldsymbol{\eta}(\cdot)$ due to only evaluating $\boldsymbol{\eta}(\mathbf{x})$ at a limited number of values of \mathbf{x} . We refer to a metamodel of this type, where predictions have the form of a probability distribution, as an *emulator*. A common choice of prior for $\boldsymbol{\eta}(\cdot)$ is a GP, in which case the emulator becomes a Bayesian version of GP regression. A detailed introduction to GP emulator methodology is given in the online toolkit provided by the Managing Uncertainty in Complex Models (MUCM) project (<http://mucm.aston.ac.uk/MUCM/MUCMToolkit>).

Early work on GP emulators [19, 20] focussed on computer models with a single scalar output, where a univariate GP is used to represent uncertainty in

the computer model output. More recently, attention has turned to emulating models with multiple outputs [21, 22, 23]. For a model with r outputs, we represent our uncertainty in the function $\boldsymbol{\eta}(\cdot)$ by the GP prior

$$\begin{aligned}\boldsymbol{\eta}(\cdot) &= \mathbf{m}(\cdot) + \mathbf{z}(\cdot), \\ \mathbf{m}(\cdot) &= (\mathbf{I} \otimes \mathbf{h}(\cdot)^T)\boldsymbol{\beta}, \\ \mathbf{z}(\cdot)|\boldsymbol{\theta} &\sim GP_r[0, C(\cdot, \cdot)].\end{aligned}\tag{2}$$

The function $\mathbf{m}(\cdot)$ is the prior mean, in which $\mathbf{h}(\cdot)$ is a vector of q regressors and $\boldsymbol{\beta}$ is a vector of rq unknown coefficients. The non-constant mean function is a type of response surface that represents the global trend of the model output across input space. This helps the emulator to predict outputs in regions of the input space which are sparsely populated by training data. We use a linear trend surface, so $\mathbf{h}(\mathbf{x})^T = (1 \ \mathbf{x}^T)$, and we assume weak knowledge about the coefficients, using the improper prior $\pi(\boldsymbol{\beta}) \propto 1$. The residual r -variate GP, $\mathbf{z}(\cdot)$, is a nonparametric component that represents the nonlinear aspects of the model response, and ensures that the emulator prediction function interpolates the training data. It has a $r \times r$ matrix-valued covariance function, $C(\cdot, \cdot)$, which is controlled by some hyperparameters $\boldsymbol{\theta}$.

3.1. Multivariate covariance functions

A key issue in constructing a multivariate GP emulator is defining a structure for $C(\cdot, \cdot)$, the covariance function of the residual process. There are two types of correlation that must be accounted for within $C(\cdot, \cdot)$. First, there are between-output correlations which arise because we expect some or all of the outputs to have a shared dependency on the physical processes within the model. An example within the context of FE models is that all modes of vibration relating to a particular component will be influenced by any uncertainty in the material properties of that component. Second, there is correlation over the input space, which arises because we expect that knowledge of the model output values at one point in the parameter space will be informative about the output values at neighbouring points.

Often in multivariate GP emulators the two types of correlation are treated as being separable, so the covariance function is the product of a between-outputs covariance matrix and a spatial correlation function (where the term ‘spatial’ in this context refers to the space of the input parameters.) Separability of the covariance structure leads to various mathematical simplifications, but may be too restrictive for applications such as FE models in which the outputs represent more than one type of physical quantity [22]. The reason is that a separable covariance has a single spatial correlation function for all the outputs, making it unsuitable for use when outputs exhibit different types of response to changes in the input parameters.

The alternative to a separable covariance is a nonseparable covariance function, which allows different spatial correlation function for each output. Methods for constructing nonseparable covariance functions are discussed in [24].

Here, we use a method based on the *linear model of coregionalization* (LMC), a tool that is popular in geostatistics for modelling multivariate spatial processes [16, 25, 26, 27]. The idea behind the LMC is to construct output processes as linear combinations of a number of building-block processes. Using the LMC, the residual process in the GP emulator prior is

$$\mathbf{z}(\cdot) = \mathbf{T}\mathbf{u}(\cdot), \quad (3)$$

where \mathbf{T} is a full-rank $r \times r$ matrix, and $\mathbf{u}(\cdot)$ is a vector of r independent zero mean GPs with unit variance and spatial correlation functions $\kappa_1(\cdot, \cdot), \dots, \kappa_r(\cdot, \cdot)$. Note that under this definition, $\mathbf{z}(\cdot)$ is still a GP, since GPs are closed under linear combination.

It follows from (3) that the covariance function for $\mathbf{z}(\cdot)$ is

$$\begin{aligned} C(\cdot, \cdot) &= \mathbf{T}[\text{diag}\{\kappa_1(\cdot, \cdot), \dots, \kappa_r(\cdot, \cdot)\}]\mathbf{T}^T \\ &= \sum_{j=1}^r \Sigma_j \kappa_j(\cdot, \cdot), \end{aligned} \quad (4)$$

where, for $j = 1, \dots, r$, $\Sigma_j = \mathbf{t}_j \mathbf{t}_j^T$, in which \mathbf{t}_j is the j th column of \mathbf{T} . The set $\{\kappa_1(\cdot, \cdot), \dots, \kappa_r(\cdot, \cdot)\}$ forms a basis of correlation functions, and the covariance function for a single output or a pair of outputs is a weighted sum of those basis functions. The weights are determined by the elements of the matrices Σ_j , $j = 1, \dots, r$, known as the *coregionalization matrices*. An advantage of the LMC construction is that, by composing the overall correlation function as combination of basis functions, variation in the response occurring on several different scales can be modelled.

3.2. Building the emulator

To build the emulator, a training design $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ is selected. A common choice for the training design is the Latin hypercube design (LHD) [28], which guarantees to spread design points evenly across each input parameter dimension. There are many different LHDs of any given dimensionality and size, so one typically generates a large number of them and chooses the one which best fulfills some space filling criterion. In this work we use LHDs chosen to maximise the minimum distance between pairs of points (the *maximin* criterion).

The computer model is run at each point in the training design, yielding a vector of training data outputs \mathbf{Y} . We condition the prior (2) on the training data and integrate over the prior distribution of the coefficients β to obtain the posterior process

$$\boldsymbol{\eta}(\cdot) | \mathbf{Y}, \boldsymbol{\theta} \sim \text{GP}_r\{\mathbf{m}^\dagger(\cdot), \mathbf{C}^\dagger(\cdot, \cdot)\}, \quad (5)$$

where, for prediction at a set of \hat{n} new input points $\hat{\mathbf{X}} = (\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_{\hat{n}})$, the poste-

rior mean and covariance are

$$\begin{aligned}\mathbf{m}^\dagger(\hat{\mathbf{X}}) &= \hat{\mathbf{H}}(\hat{\mathbf{X}})\hat{\boldsymbol{\beta}} + \mathbf{F}(\hat{\mathbf{X}})\mathbf{V}^{-1}(\mathbf{Y} - \mathbf{H}\hat{\boldsymbol{\beta}}), \\ \mathbf{C}^\dagger(\hat{\mathbf{X}}, \hat{\mathbf{X}}) &= \mathbf{C}(\hat{\mathbf{X}}, \hat{\mathbf{X}}) - \mathbf{F}(\hat{\mathbf{X}})\mathbf{V}^{-1}\mathbf{F}(\hat{\mathbf{X}})^T + \\ &\quad (\hat{\mathbf{H}}(\hat{\mathbf{X}}) - \mathbf{F}(\hat{\mathbf{X}})\mathbf{V}^{-1}\mathbf{H})(\mathbf{H}^T\mathbf{V}^{-1}\mathbf{H})^{-1}(\hat{\mathbf{H}}(\hat{\mathbf{X}}) - \mathbf{F}(\hat{\mathbf{X}})\mathbf{V}^{-1}\mathbf{H})^T,\end{aligned}$$

where $\hat{\boldsymbol{\beta}} = (\mathbf{H}^T\mathbf{V}^{-1}\mathbf{H})^{-1}\mathbf{H}^T\mathbf{V}^{-1}\mathbf{Y}$. The notation here is as follows:

$$\begin{aligned}\mathbf{H} &= \mathbf{I} \otimes \mathbf{h}(\mathbf{X})^T, \\ \hat{\mathbf{H}}(\hat{\mathbf{X}}) &= \mathbf{I} \otimes \mathbf{h}(\hat{\mathbf{X}})^T, \\ \mathbf{V} &= \mathbf{C}(\mathbf{X}, \mathbf{X}), \\ \mathbf{F}(\hat{\mathbf{X}}) &= \mathbf{C}(\hat{\mathbf{X}}, \mathbf{X}).\end{aligned}$$

The posterior mean, $\mathbf{m}^\dagger(\cdot)$, smoothly interpolates the training data and is used for point predictions of the outputs at untried input points. The uncertainty in those predictions is quantified by the posterior covariance function $\mathbf{C}^\dagger(\cdot, \cdot)$.

The GP emulator is demonstrated in Figure 1, using a toy model with one input and one output. The plot shows the posterior mean and the 95% posterior credible intervals, highlighting two important features of the emulator:

- The posterior mean, $\mathbf{m}^\dagger(\cdot)$, replicates the training data exactly.
- The posterior variance is zero at each training data point, so there is no uncertainty at points where the model has been run.

A consequence of these two features is that the emulator introduces no additional uncertainty into the analysis: the emulator merely fills the gaps between points where the model has been run with probabilistic predictions.

4. Uncertainty analysis using an emulator

The simplest way of using an emulator to carry out an uncertainty analysis is to use the posterior mean function $\mathbf{m}^\dagger(\cdot)$ as a direct replacement for $\boldsymbol{\eta}(\cdot)$ in the Monte Carlo procedure detailed in section 2. Once the training data covariance matrix \mathbf{V} has been inverted, $\mathbf{m}^\dagger(\cdot)$ can be evaluated at any input point with virtually no computational cost, so the Monte Carlo procedure becomes trivial. However, the problem with this approach is that no account is made for the extra uncertainty induced by the non-exact approximation of $\boldsymbol{\eta}(\cdot)$ by $\mathbf{m}^\dagger(\cdot)$ at untried input points.

A better approach is to note that any summary of the uncertainty distribution $S(Y)$ is itself an uncertain quantity, since it is a function of the unknown function $\boldsymbol{\eta}(\cdot)$. The following Monte Carlo procedure, proposed in [29], incorporates the uncertainty in $\boldsymbol{\eta}(\cdot)$ into the analysis:

1. Draw a sample $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_x}\}$ from the input distribution $F(\mathbf{x})$.

2. Draw a random function $\boldsymbol{\eta}_{(j)}(\cdot)$ from the emulator posterior distribution (eqn. 5) and evaluate $\boldsymbol{\eta}_{(j)}(\mathbf{x}_1), \dots, \boldsymbol{\eta}_{(j)}(\mathbf{x}_{N_x})$.
3. Obtain $S_j(Y)$, the Monte Carlo estimate of $S(Y)$ using the sample $\{\boldsymbol{\eta}_{(j)}(\mathbf{x}_1), \dots, \boldsymbol{\eta}_{(j)}(\mathbf{x}_{N_x})\}$.
4. Repeat steps 1-2 to obtain a sample $\mathcal{S}(Y) = \{S_1(Y), \dots, S_{N_\eta}(Y)\}$.
5. Use $\mathcal{S}(Y)$ to estimate any summary of the distribution of $S(Y)$.

This Monte Carlo procedure is illustrated in Figure 2, again using a toy model with one input and one output. In the example the aim is to estimate $S(Y) = \mathbb{E}[Y]$, the mean of the uncertainty distribution. Figure 2a shows the true model output as a function of the input, with the observations that are used to train the emulator. The sample from the input distribution (step 1) and the draws from the emulator posterior (step 2) are shown in Figure 2b. Each draw from the emulator posterior is evaluated at the input sample points and averaged to give an estimate of $S(Y)$ (step 3). The collection of estimates, $\mathcal{S}(Y)$, shown in Figure 2c, allows us to estimate summaries of the distribution of $S(Y)$. For example, the sample median of $\mathcal{S}(Y)$, which we denote $\hat{S}(Y)$, gives a point estimate of $S(Y)$, and credible intervals for $S(Y)$ are given by the quantiles of $\mathcal{S}(Y)$.

We note that the size of the Monte Carlo sample used in this procedure, N_x , is not limited by computational expense as before. This is because draws from the emulator posterior distribution are cheap to evaluate. As a result we can make the Monte Carlo estimate of $S(Y)$ obtained in step 3 as precise as we wish, so we consider the error to be effectively zero. In other words, the uncertainty in the summary $S(Y)$ is not due to Monte Carlo error, but rather it stems directly from our uncertainty in the function $\boldsymbol{\eta}(\cdot)$.

5. Uncertainty analysis of the FRF

We consider two approaches to analysing the uncertainty distribution of the FRF of the structure using an emulator. Recall that the random variable G_ω represents the FRF at frequency ω , and we consider obtaining arbitrary summaries $S(\cdot)$ of the distribution of G_ω .

5.1. Pointwise analysis

The most straightforward approach is to consider the FRF at a particular fixed frequency ω^* . To obtain the summary $S(G_{\omega^*})$, we use a sequential Monte Carlo procedure similar to that given above, but in step 3 we obtain the Monte Carlo estimate of $S(G_{\omega^*})$ using the sample $\{\boldsymbol{\eta}_{(j)}(\mathbf{x}_1), \dots, \boldsymbol{\eta}_{(j)}(\mathbf{x}_{N_x})\}$. We repeat steps 1-3 to obtain a sample $\mathcal{S}(G_{\omega^*}) = \{S_1(G_{\omega^*}), \dots, S_{N_\eta}(G_{\omega^*})\}$, from which we summarise the distribution $S(G_{\omega^*})$. In particular, we use the median of the sample $\mathcal{S}(G_{\omega^*})$, denoted $\hat{S}(G_{\omega^*})$, as a point estimate of $S(G_{\omega^*})$.

If we are interested in the FRF over a range of frequencies $\Omega = [\omega_1, \omega_2]$, we choose a finite number of representative values $\omega_1^*, \dots, \omega_{N_\omega}^*$ across the range

and apply the above procedure at each representative value. We refer to this approach as a pointwise analysis.

A drawback of a pointwise analysis is that a plot of summaries of $S(G_{\omega_1^*}), \dots, S(G_{\omega_{N_\omega}^*})$ may be misleading. Suppose, for example, that we choose $S(G_{\omega^*})$ to be the upper 95% quantile of G_{ω^*} . It would be natural to compute $\hat{S}(G_{\omega_1^*}), \dots, \hat{S}(G_{\omega_{N_\omega}^*})$, and to interpolate these values with a curve $\mathcal{C}_{0.95}$. It would then be tempting and to interpret $\mathcal{C}_{0.95}$ as an estimate of the upper 95% uncertainty bound on the collection of all the possible FRFs that can be obtained from the FE model. This interpretation would be wrong, though, since the FRF arising from a particular input point \mathbf{x} could lie below $\mathcal{C}_{0.95}$ at one frequency ω_i^* but above $\mathcal{C}_{0.95}$ at a different frequency ω_j^* . Thus there is no guarantee that any FRF will lie below $\mathcal{C}_{0.95}$ at all $\omega \in \Omega$.

A further drawback of a pointwise analysis is that a summary of G_{ω^*} such as the median, $S_{0.5}(G_{\omega^*})$, which we might expect to give us a central estimate of the FRF, may bear little resemblance to any actual FRF when calculated pointwise. In particular, $S_{0.5}(G_{\omega^*})$ will usually be finite for all values of ω^* , but we know that $G(\omega)$ (computed, as we do here, for an undamped structure) goes to infinity at $\omega = \sqrt{\hat{k}_i/\hat{m}_i}$ for each $i = 1, \dots, n_{modes}$. This property is demonstrated in the application in section 6.

5.2. Uncertainty envelope by kernel density estimation

An alternative to the pointwise analysis is to obtain an *uncertainty envelope*. For some fixed $\alpha \in [0, 1]$, we define an uncertainty envelope to be a set

$$E_\alpha = \{(\omega, G(\omega, \mathbf{y})) \in \mathbb{R}^2 : \omega \in \Omega, \mathbf{y} \in R\}, \quad (6)$$

where R is a subset of modal parameter space such that $Pr[\mathbf{y} \in R] = \alpha$. The uncertainty envelope has the property that the probability that G_ω lies outside of E_α for some $\omega \in \Omega$ is $1 - \alpha$.

Equation (6) does not define a unique uncertainty envelope, since there are many regions R with the property $Pr[\mathbf{y} \in R] = \alpha$. We propose choosing R to be the highest probability density (HPD) region of modal parameter space, so that for any $\mathbf{y}_1 \in R$ and $\mathbf{y}_2 \notin R$, $p_Y(\mathbf{y}_1) \geq p_Y(\mathbf{y}_2)$. We call this the *HPD α -envelope*. We note that the HPD region of modal parameter space is not necessarily connected, so the HPD α -envelope could potentially have ‘holes’ in it. To simplify the problem, we conservatively estimate the HPD α -envelope as being the region enclosed by the functions

$$\begin{aligned} \ell_\alpha(\omega) &= \min_{\mathbf{y} \in R} \{G(\omega; \mathbf{y})\}, \\ u_\alpha(\omega) &= \max_{\mathbf{y} \in R} \{G(\omega; \mathbf{y})\}. \end{aligned}$$

In other words we obtain the outer bounds of the uncertainty envelope but ignore any holes in the interior.

A difficulty in calculating the bounds of the HPD α -envelope is that we do not have an expression for the density function $p_Y(\mathbf{y})$. Instead, we estimate

$p_Y(\mathbf{y})$ using a kernel density estimate (KDE) [30] of a Monte Carlo sample. As before, we use evaluations of the emulator to provide a sequence of Monte Carlo estimates. The algorithm is as follows.

1. Choose a set of representative values $\{\omega_i \in \Omega : i = 1, \dots, N_\omega\}$.
2. Draw a sample $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_x}\}$ from the input distribution $F(\mathbf{x})$.
3. Draw a random function $\boldsymbol{\eta}_{(j)}(\cdot)$ from the emulator posterior distribution (eqn. 5) and evaluate $\mathcal{Y}_j = \{\mathbf{y}_1^{(j)}, \dots, \mathbf{y}_{N_x}^{(j)}\}$, where for $i = 1, \dots, N_x$, $\mathbf{y}_i^{(j)} = \boldsymbol{\eta}_{(j)}(\mathbf{x}_i)$.
4. Obtain $\hat{p}_Y^{(j)}(\cdot)$, the KDE of the sample \mathcal{Y}_j .
5. Let $\mathcal{P}_j = \{\mathbf{y}^{(j)} \in \mathcal{Y}_j : \hat{p}_Y^{(j)}(\mathbf{y}^{(j)}) \geq p^*\}$, where p^* is the largest number such that $\sum_{i=1}^{N_x} I[\hat{p}_Y^{(j)}(\mathbf{y}_i^{(j)}) \geq p^*] > N\alpha$.
6. Let $\ell^{(j)}$ be the linear interpolation of the set $\{(\omega_i, \min_{\mathbf{y} \in \mathcal{P}_j} \{G(\omega_i; \mathbf{y})\}) : i = 1, \dots, N_x\}$, and let $u^{(j)}$ be the linear interpolation of the set $\{(\omega_i, \max_{\mathbf{y} \in \mathcal{P}_j} \{G(\omega_i; \mathbf{y})\}) : i = 1, \dots, N_x\}$.
7. Let $\hat{E}_\alpha^{(j)}$ be the region of $\Omega \times \mathbb{R}$ enclosed by $\ell^{(j)}$ and $u^{(j)}$.
8. Repeat steps 3-7 to obtain a sample $\mathcal{E}_\alpha = \{\hat{E}_\alpha^{(1)}, \dots, \hat{E}_\alpha^{(N_\eta)}\}$

We summarise the sample \mathcal{E}_α to obtain an estimate and credible intervals for the uncertainty envelope E_α . Note that if we choose $\alpha = 0$ then this algorithm returns an estimate of the FRF corresponding to the mode of the uncertainty distribution, as it should.

The HPD α -envelope algorithm is illustrated in Figure 3. In the illustration there are two inputs $\mathbf{x} = (x_1, x_2)^T$ with a uniform distribution, from which a sample is taken (step 2; Figure 3a). This sample is mapped into modal parameter space by a draw from the emulator posterior (step 3). The KDE of the sample in modal parameter space is used to rank the points according to their probability density, and the top $100\alpha\%$ are selected (steps 4-5; Figure 3b). The FRF is computed for each selected point, and the collection of FRFs is used to construct the envelope (steps 6-7; Figure 3c). The procedure is repeated with different draws from the emulator posterior, and each draw results in a different envelope. The collection of envelopes shows us the uncertainty in E_α induced by only running the FE model at a limited number of input configurations.

A disadvantage of the uncertainty envelope approach is that, given a pointwise analysis, carried out at the α level and at $N_\omega > 1$ discrete frequencies, E_α will be wider than at least one of the pointwise intervals. This is because the probability of the FRF lying in every pointwise uncertainty interval simultaneously is less than α . Therefore, if interest is only in the FRF at individual frequencies then the pointwise analysis may be preferable. If, however, interest is in the location of the entire FRF within a given range of frequencies, then the uncertainty envelope approach should be taken.

Fuselage length (mm)	1500
Wingspan (mm)	2000
Fuselage thickness (mm)	50
Tail tip mass \mathbf{x}_1 (kg)	0.4 - 0.6
Wingtip mass \mathbf{x}_2 (kg)	0.16 - 0.24
Wing trailing edge thickness \mathbf{x}_3 (mm)	10.4 - 11.6
Wing leading edge / end plate thickness \mathbf{x}_4 (mm)	9.5 - 10.5
Young's modulus \mathbf{x}_5 (GPa)	64.8 - 79.2
Density (kg/m ³)	2700
Poisson's ratio	0.34

Table 1: Parameters used for the Garteur SM-AG-19 testbed FE model

6. Application

The aircraft model of the Garteur benchmark problem was designed by the Garteur Structures and Materials Action Group 19 (SM-AG-19) to evaluate ground vibration test techniques [31]. The testbed has since been used for a variety of benchmarking and case study problems in structural dynamics, such as FE model updating [32], and fuzzy FE [6]. In the present study, an FE representation of the testbed is used to explore the effectiveness of the proposed statistical emulator approach to uncertainty analysis.

The FE model is based upon the *gartfe* model included in the Matlab Structural Dynamics Toolbox (www.sdtools.com). In the physical testbed, a constrained layer damping treatment is applied to the trailing edge of the wings to increase the structural damping. However, for the present study the structural damping is neglected so as to focus on the undamped FRF predictions that arise directly from the modes of vibration that can be calculated from the FE model. Meanwhile, five parameters of the FE model, denoted $\mathbf{x}^T = (x_1, \dots, x_5)$, are chosen to be uncertain, with a uniform distribution over a hypercuboidal input space. The resulting physical parameters of the model are listed in Table 1. In practice, such uncertainties may arise due to unknown parameters during the iterative design process, manufacturing inaccuracies/variabilities, or changes in operating conditions. However, for the purposes of the present study the choice of which parameters are uncertain (and their distributions) is somewhat arbitrary.

The model structure and resulting modes of vibration are illustrated in Figure 4. In order to consider the FRF of the structure, a single degree of freedom of the FE model is selected as shown in 4a, and the effective modal mass (kg) and stiffness (N/m) obtained for each of the first 3 structural modes of vibration. These are referred to as $(\hat{m}_1, \hat{k}_1, \hat{m}_2, \hat{k}_2, \hat{m}_3, \hat{k}_3)$. The effective mass of the rigid-body modes of vibration, \hat{m}_{rig} , was also obtained in order to obtain mass-dominated receptance FRFs as would be expected for the freely-suspended structure. The combined output vector is $\mathbf{y}^T = (\hat{m}_1, \hat{k}_1, \hat{m}_2, \hat{k}_2, \hat{m}_3, \hat{k}_3, \hat{m}_{rig}) = (y_1, \dots, y_7)$

6.1. Building and validating the emulator

We compare four metamodels for the FE model:

- GP_{MV} : A multivariate GP emulator using the LMC covariance function.
- GP_{ind} : A collection of seven independent univariate GP emulators, one for each output.
- RS_{lin} : A collection of seven independent linear response surfaces. Each linear RS is a first degree polynomial with pairwise interactions, of the form $y = \alpha_0 + \sum_{i=1}^5 \alpha_i x_i + \sum_{j>i} \beta_{ij} x_i x_j$.
- RS_{quad} : A collection of seven independent quadratic response surfaces. Each quadratic RS is a second degree polynomial with pairwise first degree interactions of the form $y = \alpha_0 + \sum_{i=1}^5 \alpha_i x_i + \sum_{j>i} \beta_{ij} x_i x_j$.

To train the metamodels we use $n = 50$ runs of the FE model in a Latin hypercube design that covers the input space. We also have available a further 20 runs of the FE model, distinct from the training data, that we use to test our emulator. We predict the outputs of the test data using the metamodels.

We find that the predictions of \hat{m}_{rig} have close to zero error for all four metamodels. This is because \hat{m}_{rig} is a very smooth and almost linear function of the input parameters that correspond to physical masses and dimensions (i.e. parameters \mathbf{x}_1 to \mathbf{x}_4).

To assess the accuracy of the metamodels in predicting the other outputs we compute the standardised RMS prediction error for each structural mode of vibration as

$$RMSE_j = \sqrt{\sum_{i=1}^{20} \left(\frac{e_{\hat{m}_j}^{(i)}}{r_{\hat{m}_j}} \right)^2 + \left(\frac{e_{\hat{k}_j}^{(i)}}{r_{\hat{k}_j}} \right)^2}, \quad (7)$$

where $e_{\hat{m}_j}^{(i)}$ and $e_{\hat{k}_j}^{(i)}$ are the prediction errors for modal parameters \hat{m}_j and \hat{k}_j respectively, and $r_{\hat{m}_j}$ and $r_{\hat{k}_j}$ are the ranges of the validation data values for \hat{m}_j and \hat{k}_j respectively. The standardised RMS prediction errors are compared in Figure 5, and the relative accuracies of the metamodels are illustrated further in Figure 6, where we plot validation data and predictions in (\hat{m}_i, \hat{k}_i) cross-sections of the output space. We see that the prediction accuracies of GP_{MV} and GP_{IND} are almost identical for all three modes. For modes 1 and 2, RS_{quad} and the GP emulators have very similar RMSEs, which are relatively small compared to the RMSE of RS_{lin} . For mode 3 RS_{quad} has worse predictions than RS_{lin} , while the GP emulators do better than both the RS metamodels. This suggests that the response of the FE model is non-linear for all modes, and non-quadratic for mode 3, so the response surfaces we have chosen are not adequate. RS_{quad} does worse than RS_{ind} because, with 31 polynomial coefficients to estimate from just 50 training data, it is somewhat overfitted and likely to predict very poorly in regions that deviate from quadratic behavior. The GP emulators

are nonparametric, so are better able to adapt to non-linear models than the response surfaces.

Figure 7 shows (\hat{m}_i, \hat{k}_i) cross-sections of the output space, zoomed in on a single validation point. Also shown in these plots are 95% highest posterior density credible regions for the GP emulator predictions, calculated marginally for each pair of outputs (\hat{m}_i, \hat{k}_i) . These credible regions are ellipses that represent the joint predictive uncertainty in each pair of modal parameters. Here we see that the main difference between the two GP emulators is in the shape of the credible regions, most notably in mode 4. GP_{IND} considers \hat{m}_i and \hat{k}_i to be independent, so its credible regions are ellipses orientated with the coordinate axes. GP_{MV} , on the other hand, accounts for dependency between \hat{m}_i and \hat{k}_i , so its credible regions are ellipses orientated in the direction of greatest uncertainty. Consequently, the credible regions of GP_{MV} better reflect the joint uncertainty in the modal parameters than do the credible regions of GP_{IND} . The advantage of this will be seen subsequently.

We have demonstrated here that the GP emulators have overall better prediction accuracy than the linear and quadratic RS metamodels. They also have the advantage of quantifying the uncertainty in their predictions. We therefore drop the RS metamodels, using just the GP emulators for the subsequent analyses.

6.2. Pointwise uncertainty analysis

We carry out a pointwise uncertainty analysis of the FRF using the method detailed in section 5.1 with each of the GP emulators in turn. Figure 8 shows the estimates of the median and the 50% and 95% equal-tail uncertainty intervals, calculated pointwise for 100 values of ω spanning the range $[0, 300]$. We see that the 95% credible intervals for the median and the bounds on the uncertainty intervals are wider with GP_{IND} than with GP_{MV} , showing that using a multivariate GP rather than independent univariate GPs results in less uncertainty in the estimate of the distribution of G_ω . The reason for this is that when the model parameters are combined in the calculation of the FRF, the lack of covariance structure in the independent emulators results in an artificial inflation of the FRF variance.

The drawbacks of a pointwise analysis that we discuss in section 5.1 are shown in Figure 9, where we impose the FRFs arising from all 70 evaluations of the FE model that we have available (i.e. both the training and validation data) onto the 95% pointwise equal-tail uncertainty intervals. A total of 63 of these FRFs fall outside of the interval for some value of ω , demonstrating that the pointwise analysis can give a misleading representation of the uncertainty in the FRF. Also, the pointwise median of the uncertainty distribution in Figures 8e-8f bears little resemblance to any individual FRF. In particular, the pointwise median is finite everywhere, so we cannot use it for point predictions of the undamped natural frequencies.

6.3. Uncertainty envelope analysis

Figure 10 shows estimated HPD uncertainty envelopes from each of the two GP emulators at the levels $\alpha = 0, 0.5, 0.95$. (Recall that the HPD uncertainty envelope at $\alpha = 0$ corresponds to the FRF at the mode of the uncertainty distribution.) The 95% credible intervals for bounds on the uncertainty envelopes are considerably wider with GP_{IND} than with GP_{MV} , again because ignoring the covariance structure results in inflation of the FRF variance. Indeed, the credible intervals with GP_{IND} for the 95% uncertainty envelope are so wide that we would probably conclude that further evaluations of the FE model are required in order to reduce the emulator induced uncertainty to an acceptable level. With GP_{MV} , however, the same credible intervals are much smaller and the extra evaluations may not be required.

As we expect, the HPD uncertainty envelopes are wider than the interpolated pointwise intervals corresponding to the same probability levels. In Figure 11 we see that all of the FRFs arising from the emulator training and validation data are fully contained within the 95% HPD uncertainty envelope, demonstrating that the uncertainty envelope gives a better overall picture of the region in which we expect to find the FRF.

7. Conclusions

The contribution of this paper is twofold. First, a probabilistic method for quantifying uncertainty in an FRF of a structure is proposed. It is well known that a pointwise analysis (in which an independent uncertainty interval for the FRF is calculated at a finite number of frequencies) can give misleading information about the FRF as whole. The present study has developed an alternative uncertainty envelope approach, in which an envelope function corresponding to the region of highest probability density in the modal parameter space is calculated. The uncertainty envelope tends to be wider than the corresponding set of pointwise intervals, but has the advantage that it gives a true representation of uncertainty about the location of the entire continuous FRF within a given range of frequencies.

The second contribution of this paper is to introduce the multivariate GP emulator as a tool for performing such probabilistic uncertainty analyses. In the early stages of development of the SFE method it was considered adequate to compute low-order statistics of the variables of interest, but uncertainty analysis for FE methods has become increasingly sophisticated in recent times. The current state-of-the-art requires estimates of more complex quantities such as measures of reliability, which often can only be computed via Monte Carlo methods. The GP emulator approach introduced in this paper is ideally suited to such applications, since traditional Monte Carlo methods are usually prohibitively expensive in large FE models. The GP emulator has two distinct advantages over other types of metamodel:

- The GP emulator is a form of nonparametric regression that exactly interpolates the training data, so there is no error or uncertainty about the

output at input points where the FE model has been run.

- The Bayesian approach means that the uncertainty about the output at new input points is quantified and propagated through the rest of the analysis, giving credible intervals for the bounds on the pointwise intervals and the uncertainty envelope.

In a practical example we have demonstrated that GP emulators can give more accurate predictions than linear and quadratic response surfaces. We have also demonstrated that a multivariate GP emulator is better at quantifying uncertainty in multiple outputs than a collection of independent univariate GP emulators. This is because there are dependencies between the outputs, and modelling the uncertainties in them jointly results in a significant reduction in the uncertainty that is propagated through the subsequent analyses. As such, we have shown the multivariate GP emulator to be a very useful tool in the study of multiple output FE models. Further work is needed to extend this approach to consider the problems of overlapping modes of vibration, and uncertain structural damping.

Acknowledgements

This research is part of the ‘Managing Uncertainty in Complex Models’ (MUCM) project, funded by Research Councils UK under its Basic Technology programme, and received support from the EPSRC grants ‘Uncertainty Propagation in Structures, Systems and Processes’ (EP/D078601/1) and ‘Bridging the Gaps’ (EP/E01867X/1). The authors thank the other members of the MUCM team for helpful comments.

References

- [1] R. Ghanem, P. Spanos, *Stochastic finite elements: a spectral approach*, Dover Publications, Mineola N.Y., 2003.
- [2] B. Lallemand, A. Cherki, T. Tison, P. Level, Fuzzy modal finite element analysis of structures with imprecise material properties, *Journal of Sound and Vibration* 220 (1999) 353–365.
- [3] G. Plessis, B. Lallemand, T. Tison, P. Level, Fuzzy modal parameters, *Journal of sound and vibration* 233 (2000) 797–812.
- [4] L. Hinke, F. Dohnal, B. Mace, T. Waters, N. Ferguson, Component mode synthesis as a framework for uncertainty analysis, *Journal of Sound and Vibration* 324 (2009) 161–178.
- [5] S. Donders, D. Vandepitte, J. Van de Peer, W. Desmet, Assessment of uncertainty on structural dynamic responses with the short transformation method, *Journal of Sound and Vibration* 288 (2005) 523–549.

- [6] H. De Gerssem, D. Moens, W. Desmet, D. Vandepitte, A fuzzy finite element procedure for the calculation of uncertain frequency response functions of damped structures: Part 2–Numerical case studies, *Journal of sound and vibration* 288 (2005) 463–486.
- [7] A. O’Hagan, C. Buck, A. Daneshkhah, J. Eiser, P. Garthwaite, D. Jenkinson, J. Oakley, T. Rakow, *Uncertain judgements: eliciting experts’ probabilities*, John Wiley, Chichester, 2006.
- [8] J. Hammersley, D. Handscomb, *Monte carlo methods*, Taylor & Francis, 1964.
- [9] E. El Tabach, L. Lancelot, I. Shahrouh, Y. Najjar, Use of artificial neural network simulation metamodelling to assess groundwater contamination in a road project, *Mathematical and Computer Modelling* 45 (2007) 766–776.
- [10] H. Drucker, C. Burges, L. Kaufman, A. Smola, V. Vapnik, Support vector regression machines, *Advances in neural information processing systems* (1997) 155–161.
- [11] L. Wang, S. GmbH, N. Springer, *Support Vector Machines: theory and applications* (2005).
- [12] T. Lew, A. Spencer, F. Scarpa, K. Worden, A. Rutherford, F. Hemez, Identification of response surface models using genetic programming, *Mechanical Systems and Signal Processing* 20 (2006) 1819–1831.
- [13] J. Dejean, G. Blanc, Managing uncertainties on production predictions using integrated statistical methods, in: *Society of Petroleum Engineers Annual Technical Conference and Exhibition*.
- [14] A. Rutherford, D. Inman, G. Park, F. Hemez, Use of response surface meta-models for identification of stiffness and damping coefficients in a simple dynamic system, *Shock and Vibration* 12 (2005) 317–331.
- [15] A. O’Hagan, Curve fitting and optimal design for prediction, *Journal of the Royal Statistical Society. Series B* 40 (1978) 1–42.
- [16] A. G. Journel, C. J. Huijbregts, *Mining Geostatistics*, Academic Press, 1978.
- [17] N. Cressie, *Statistics for Spacial Data*, John Wiley, 1991.
- [18] C. E. Rasmussen, C. K. I. Williams, *Gaussian Processes for Machine Learning*, MIT Press, 2006.
- [19] C. Currin, T. Mitchell, M. Morris, D. Ylvisaker, Bayesian prediction of deterministic functions, with applications to the design and analysis of computer experiments, *Journal of the American Statistical Association* 86 (1991) 953–963.

- [20] M. Morris, T. Mitchell, D. Ylvisaker, Bayesian design and analysis of computer experiments: use of derivatives in surface prediction, *Technometrics* (1993) 243–255.
- [21] S. Conti, A. O’Hagan, Bayesian emulation of complex multi-output and dynamic computer models, *Journal of statistical planning and inference* 140 (2010) 640–651.
- [22] J. Rougier, Efficient emulators for multivariate deterministic functions, *Journal of Computational and Graphical Statistics* 17 (2008) 827–843.
- [23] M. Kennedy, C. Anderson, A. O’Hagan, M. Lomas, I. Woodward, J. Gosling, A. Heinemeyer, Quantifying uncertainty in the biospheric carbon flux for England and Wales, *Journal of the Royal Statistical Society: Series A(Statistics in Society)* 171 (2008) 109–135.
- [24] T. E. Fricker, J. Oakley, N. M. Urban, Multivariate Emulators with Non-separable Covariance Structures (2010). MUCM Technical Report.
- [25] H. Wackernagel, *Multivariate Geostatistics*, Springer, 1995.
- [26] M. Goulard, M. Voltz, Linear coregionalization model: Tools for estimation and choice of cross-variogram matrix, *Journal Mathematical Geology* 21 (1992) 269–286.
- [27] A. E. Gelfand, A. M. Schmidt, S. Banerjee, C. F. Sirmans, Nonstationary multivariate process modelling through spatially varying coregionalization, *Test* 13 (2004) 1–50.
- [28] M. D. McKay, R. J. Beckman, W. J. Conover, A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics* 21 (1979) 239–245.
- [29] J. E. Oakley, A. O’Hagan, Bayesian inference for the uncertainty distribution of computer model outputs, *Biometrika* 89 (2002) 769–784.
- [30] B. Silverman, *Density estimation for statistics and data analysis*, Chapman & Hall/CRC, 1998.
- [31] E. Balmes, J. Wright, GARTEUR group on ground vibration testing: results from the test of a single structure by 12 laboratories in Europe, in: *Proceedings of the International Modal Analysis Conference*, Orlando.
- [32] C. Mares, M. Friswell, J. Mottershead, Model updating using robust estimation, *Mechanical Systems and Signal Processing* 16 (2002) 169–183.




Figure 1: Illustration of an emulator. Depicted are the true function (solid line), the training data (bullets), the posterior mean (dotted line passing through the data points) and the 95% posterior credible intervals (grey regions enclosed by dashed lines).




Figure 2: Illustration of the Monte Carlo procedure for uncertainty analysis using an emulator. (a) True function with emulator training data. (b) Sample from the input distribution (crosses and dotted grey vertical lines) and four draws from the emulator posterior (dashed black lines). (c) Estimates of $S(Y) = \mathbb{E}[Y]$, calculated using each of the realisations (plotting symbols correspond to those in (b)). The true value of $S(Y) = \mathbb{E}[Y]$ is shown by the solid circle and dashed horizontal line.




Figure 3: Illustration of the HPD α -envelope algorithm. (a) Sample from the input space distribution. (b) Sample mapped to modal parameter space using a single draw from the emulator posterior. Black dots are the points identified as having the highest probability density. (c) FRFs corresponding to each of the black points in (b) (solid lines), with the resulting uncertainty envelope (dashed line).




Figure 4: FE model of the Garteur SM-AG-19 Testbed. (a) Undeformed model, showing the degree-of-freedom (DOF) used for the frequency response function. (b-d) First three modes of structural vibration.




Figure 5: Standardised RMS prediction errors for each structural mode of vibration.

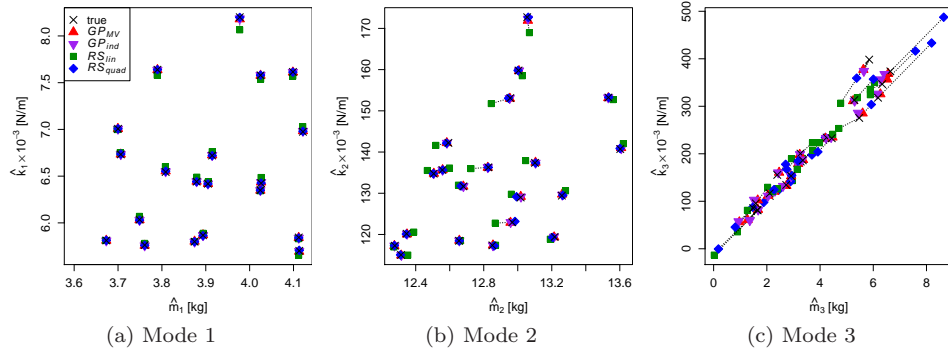


Figure 6: Predictions of validation points. Depicted are (\hat{m}_i, \hat{k}_i) cross-sections of the output space, comparing the four metamodels.

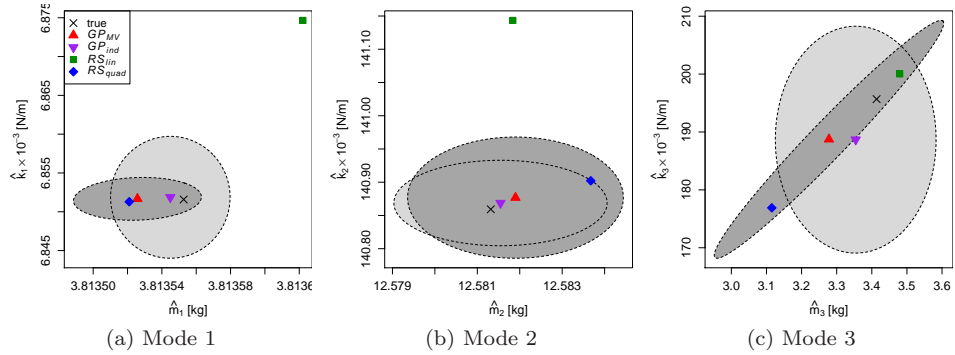


Figure 7: Prediction of a typical validation point. Depicted are (\hat{m}_i, \hat{k}_i) cross-sections of the output space, comparing the four metamodels. Light grey regions: 95% credible regions for GP_{IND} . Dark grey regions: 95% credible regions for GP_{MV} .

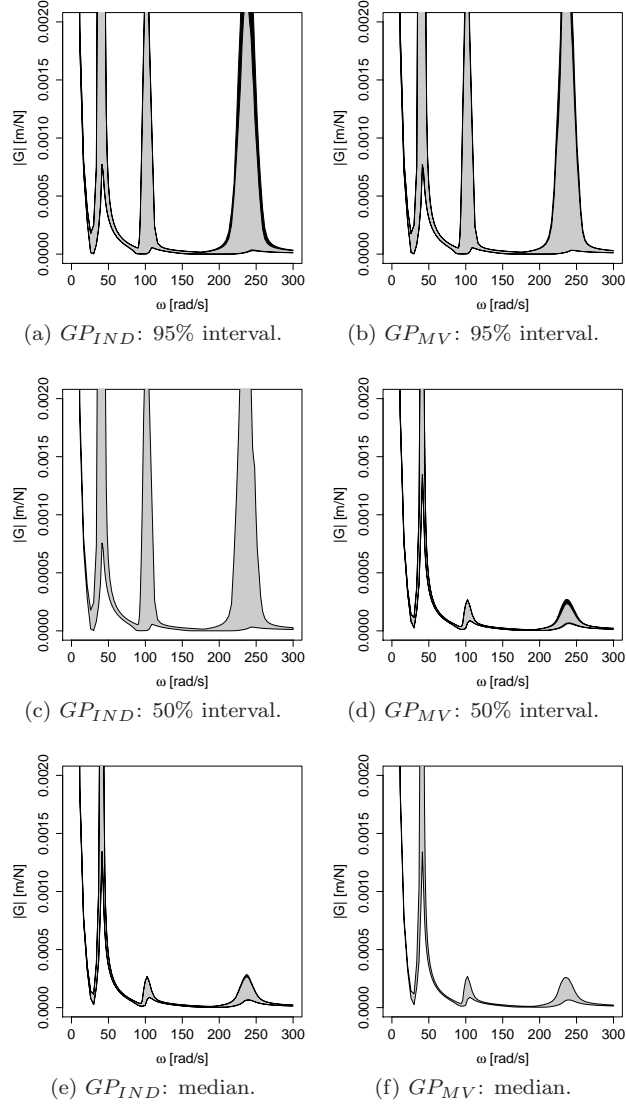


Figure 8: Pointwise analysis. Left panels correspond to GP_{IND} ; right panels correspond to GP_{MV} . (a)-(d): 95% and 50% pointwise uncertainty intervals (grey regions) and 95% credible intervals for the bounds on the pointwise uncertainty intervals (black regions). (e)-(f): 95% credible intervals for the pointwise median (black regions).




Figure 9: Pointwise analysis with data. 95% pointwise uncertainty intervals as in Figures 8a and 8b, with the 50 training data and 20 validation imposed behind (dashed lines).

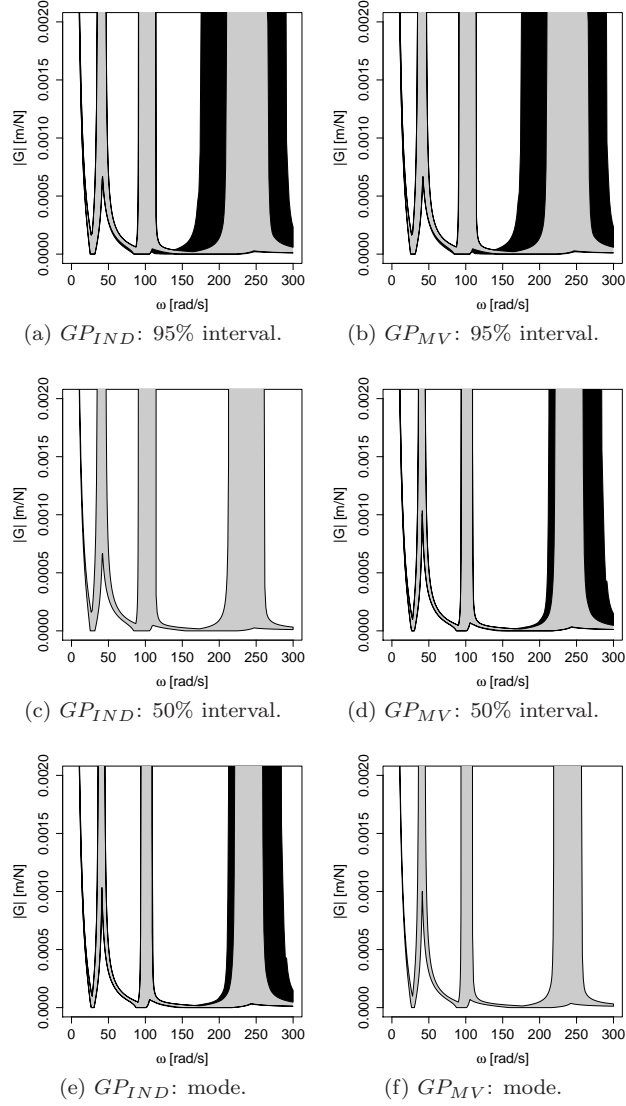


Figure 10: Uncertainty envelope analysis. Left panels correspond to GP_{IND} ; right panels correspond to GP_{MV} . (a)-(d): Estimated HPD uncertainty envelopes $E_{0.95}$ and $E_{0.5}$ (grey regions) and 95% credible intervals for the bounds on those envelopes (black regions). (a)-(b): 95% credible intervals for the estimated maximum probability FRF, E_0 (black regions).




Figure 11: Uncertainty envelope analysis with data. 95% HPD uncertainty envelopes as in Figures 10a and 10b, with the 50 training data and 20 validation superimposed (dashed lines).