



This is a repository copy of *Modelling variability in vibration-based PBSHM via a generalised population form*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/208309/>

Version: Published Version

Article:

Dardeno, T.A., Bull, L.A., Mills, R.S. orcid.org/0000-0003-1613-2369 et al. (2 more authors) (2022) Modelling variability in vibration-based PBSHM via a generalised population form. *Journal of Sound and Vibration*, 538. 117227. ISSN 0022-460X

<https://doi.org/10.1016/j.jsv.2022.117227>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of Sound and Vibration

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

Modelling variability in vibration-based PBSHM via a generalised population form

T.A. Dardeno ^{a,*}, L.A. Bull ^b, R.S. Mills ^a, N. Dervilis ^a, K. Worden ^a^a Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Sheffield S1 3JD, UK^b The Alan Turing Institute, The British Library, London NW1 2DB, UK

ARTICLE INFO

Keywords:

Repeatability

Uncertainty

Gaussian processes

Population-based SHM

ABSTRACT

Structural health monitoring (SHM) has been an active research area for the last three decades, and has accumulated a number of critical advances over that period, as can be seen in the literature. However, SHM is still facing challenges because of the paucity of damage-state data, operational and environmental fluctuations, repeatability issues, and changes in boundary conditions. These issues present as inconsistencies in the captured features and can have a huge impact on the practical implementation, but more critically, on the generalisation of the technology. Population-based SHM has been designed to address some of these concerns by modelling and transferring missing information using data collected from groups of similar structures.

In this work, vibration data were collected from four healthy, nominally-identical, full-scale composite helicopter blades. Manufacturing differences (e.g., slight differences in geometry and/or material properties), among the blades presented as variability in their structural dynamics, which can be very problematic for SHM based on machine learning from vibration data. This work aims to address this variability by defining a general model for the frequency response functions of the blades, called a *form*, using mixtures of Gaussian processes.

1. Introduction

Engineering structures are inherently uncertain, even when structures can be considered *nominally-identical*, because of manufacturing differences causing variation in geometry and material properties. These differences, as well as variations caused by ageing parts and changes in testing conditions (e.g., operational and environmental fluctuations or changes in boundary stiffness), make generalisation among structures difficult, particularly with respect to their dynamic properties [1,2]. The current work is focussed on applying structural health monitoring (SHM), to a set of nominally-identical (i.e., *homogeneous* [3,4]) structures for the purpose of damage detection. (As opposed to a *heterogeneous* population [5,6], where the members are more disparate, such as different designs of a suspension bridge). However, these fluctuations unrelated to damage are challenging for SHM, as damage-sensitive features may also be sensitive to normal variations, which may cause issues with discerning between damaged and healthy states [7–10]. For example, it is well-known that structural damage is often associated with a reduction in stiffness. Depending on the location and severity of the damage, and the complexity of the examined mode, this damage may present as a decrease in natural frequency. Increased ambient temperature, loosening of bolts at a support, or similar situations may also reduce stiffness and can mimic damage. Likewise, conditions that increase stiffness (e.g., decreased ambient temperature) can increase natural frequency and mask damage.

* Corresponding author.

E-mail address: t.a.dardeno@sheffield.ac.uk (T.A. Dardeno).<https://doi.org/10.1016/j.jsv.2022.117227>

Received 21 March 2022; Received in revised form 28 July 2022; Accepted 2 August 2022

Available online 6 August 2022

0022-460X/© 2022 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

This sensitivity of natural frequency to environmental/operational changes in addition to damage has been demonstrated experimentally [8–10]. Alampalli [8], performed impact modal testing on a small 6.76×5.26 metre bridge, before and after damaging the bridge with a saw cut made across the bottom flanges of both the main girders. They found that for temperatures above freezing, the stiffness reduction caused by damage was reflected as a downward shift in natural frequencies (–2% to –11%, depending on the mode), as expected. However, they also found that for temperatures below freezing, the measured frequencies increased significantly, primarily from freezing of accumulated moisture within the supports (+25% to +58%, depending on the mode), which masked the reduction caused by damage. In addition, uncertainty among nominally-identical structures (e.g., slight variation in internal structure and material properties), like those caused by environmental and other fluctuations, may be erroneously flagged as damage or mask subtle damage in a structure. Cawley [9], initiated a crack at the fixed end of a cantilever beam and evaluated the natural frequencies as the beam length varied. He found that the change in natural frequency resulting from a cut through 2% of the beam depth was 40 times smaller than that caused by a 2% increase in the beam length [9,10]. Given that these small differences among structures and testing conditions can have such a large impact on dynamics and damage detection, one would ideally aim to collect new data for each structure under each set of testing conditions. However, collecting comprehensive data covering a range of operational and damaged conditions from each individual structure is typically not feasible (a wind farm, for example, may contain hundreds or even thousands of wind turbines). In such cases, population-based SHM (PBSHM) can be used to make inferences between different members of a *population*, i.e., groups of similar systems [3–6,11].

PBSHM seeks to transfer valuable information across similar structures. In certain cases, the behaviour of the group can be represented using a general model. Bull et al. used Gaussian process (GP) regression of real [3,4], and imaginary [4], frequency response functions (FRFs) to develop a generic representation, called a *form*, of a population of nominally-identical, but slightly different, eight degree-of-freedom (DOF) systems. The population was comprised of multiple simulated healthy and damaged systems [3,4], and one experimental rig, which was used to test the form in [3]. Normal-condition data from several members of the population were used to train the form, and the posterior predictive distributions of the GP were used to aid novelty detection across the population, using the Mahalanobis squared-distance novelty index [12]. The form was found to predict accurately the healthy and damaged states across the population [3,4]. Lázaro-Gredilla et al. [13], introduced the overlapping mixture of Gaussian processes (OMGP), which uses a variational Bayesian approach to learn hyperparameters and labels by clustering data according to various trajectories across the input space. This technique was applied to the population form in [3,11], using an OMGP to infer multivalued wind-turbine power-curve data, with categorisation of the data performed in an unsupervised manner. Novelty detection was performed using the negative log-marginal-likelihood of the OMGP, and correctly identified *good* and *bad* power curves for the majority of data tested [3,11].

In this work, data were collected from four healthy, nominally-identical helicopter blades. FRFs were computed from the measured time-domain data, and these data were used to develop a form for the blades, first using a supervised mixture of GPs, and second using an (unsupervised) OMGP. The posterior predictive distributions from the OMGP were later applied to a novelty detection approach to evaluate new (simulated) FRF data against the form. As such, this work builds upon that performed in [3,4], by extending the concept of a population form to a set of nominally-identical helicopter blades with a range of undamaged-condition FRF data, using mixtures of probabilistic regression models [3,11,13], to account for the large variability in the FRFs of the blades. The significant discrepancies in the experimental data set are representative of the differences in the blades in practice, with variability arising from small variations in the internal structure and material properties of the helicopter blades, as well as slight changes in boundary conditions. The small population of helicopter blades provided a useful data set to develop and test the techniques presented herein, such that these methods can be later applied to other homogeneous populations such as a wind farm.

Although the specific application of this work relates to population-based damage detection, this work is also applicable to the more general problem of dynamic analysis of a population. Attempts to quantify uncertainty in vibration data have included parametric [14], and nonparametric [2,15], approaches. Parametric models allow incorporation of engineering knowledge, but can suffer from over-training (if highly complex), or over-simplification. Nonparametric models, such as the approach used in [2], and the polynomial-chaos expansion used in [15], in general do not incorporate prior engineering knowledge, and this can lead to problematic solutions because of the high model flexibility. Gaussian processes provide a good middle ground, as they are nonparametric, but allow for the use of a parametric mean function where prior knowledge about the data can be incorporated into the model. The (OM)GP form provides a probabilistic model of the FRF that can capture uncertainty caused by the variations described above as well as measurement noise.

This work formalises the use of a modal analysis approach to represent a group of similar *dynamic* systems, using a data-driven model (or representation) referred to as the *population form* [3,4]. Unlike previous work [4], the form here uses a mixture of regressions to represent *population dynamics/FRFs* — rather than a single, averaged model (complete pooling). In turn, the modelling procedure can approximate population data with more complex variations and different modes in the posterior distribution; for example, relating to distinct operating conditions, characteristic subgroups in the population, or individual members. While mixture models have been used to represent the *form* in previous work, relating to wind farm power prediction [3,11], these models were not physics informed [16]; here, in contrast, the proposed (semi-parametric) form is constructed around the established theory of linear modal analysis — allowing for the natural inclusion of domain (engineering) expertise via interpretable parameters (e.g., natural frequencies, damping ratios, mode shapes).

The layout of this paper is as follows. Section 2 provides an overview of the theory used in this work, including modal analysis, conventional GPs, and OMGP. Section 3 describes the experimental campaign performed on the population of nominally-identical helicopter blades and introduces the data set used in the analyses. Section 4 discusses how a generalised normal condition was determined for the helicopter blades, first using a supervised mixture of GPs, and second using an (unsupervised) OMGP. Finally, Section 5 demonstrates how the OMGP can be used to inform damage detection, by computing the marginal likelihood for a series of (simulated) damaged FRFs.

2. Background theory

In this work, a generalised normal condition was determined for a small population of nominally-identical helicopter blades. A form was developed for the population, first using a supervised mixture of GPs, and second using an (unsupervised) OMGP. A modal-based FRF estimation technique was parameterised and used to compute a non-zero prior mean function for the GPs. The marginal likelihood (evidence) was then used along with the posterior predictive distributions of the OMGP to evaluate (simulated) new data for novelty. A theoretical background for these methodologies is provided below.

2.1. Modal analysis

The equation of motion for a multiple DOF system can be written as,

$$\mathbf{M}\ddot{\mathbf{u}}(t) + \mathbf{C}\dot{\mathbf{u}}(t) + \mathbf{K}\mathbf{u}(t) = \mathbf{z}(t) \quad (1)$$

where $\ddot{\mathbf{u}}(t)$, $\dot{\mathbf{u}}(t)$, $\mathbf{u}(t)$, and $\mathbf{z}(t)$ are acceleration, velocity, displacement, and force, respectively. In most cases, the mass \mathbf{M} , damping \mathbf{C} , and stiffness \mathbf{K} matrices are coupled. For linear systems, and in the absence of viscous damping, the equation of motion can be decoupled, such that the system is represented by multiple single degree-of-freedom (SDOF) oscillators. This decoupling is performed via the eigenvalue expression,

$$[\mathbf{K} - \omega_n^2 \mathbf{M}] \Psi = 0 \quad (2)$$

and yields the natural frequencies in radians, ω_n , and mode shapes, Ψ , of the system.

The physical equation of motion can then be cast in modal space to give the uncoupled modal equations, with modal coordinates $\mathbf{p}(t)$, written as,

$$\Psi^T \mathbf{M} \Psi \ddot{\mathbf{p}}(t) + \Psi^T \mathbf{C} \Psi \dot{\mathbf{p}}(t) + \Psi^T \mathbf{K} \Psi \mathbf{p}(t) = \Psi^T \mathbf{z}(t) \quad (3)$$

where,

$$\mathbf{u}(t) = \Psi \mathbf{p}(t) \quad (4)$$

Frequency response functions (FRFs) can be used to visualise the natural frequency components of a system, and are computed by normalising the response signal at a given location to the excitation force. This work used the H_1 estimator to compute FRFs. For a response at location i resulting from excitation at j , the H_1 estimator is computed as,

$$\mathbf{H}_{ij}(\omega) = \frac{\mathbf{G}_{zu}(\omega)}{\mathbf{G}_{zz}(\omega)} \quad (5)$$

where,

$$\mathbf{G}_{zu}(\omega) \stackrel{\Delta}{=} \mathbb{E}[\mathbf{S}_{u_i}(\omega) \mathbf{S}_{z_j}^*(\omega)] \quad (6)$$

$$\mathbf{G}_{zz}(\omega) \stackrel{\Delta}{=} \mathbb{E}[\mathbf{S}_{z_j}(\omega) \mathbf{S}_{z_j}^*(\omega)]$$

and,

$$\mathbf{S}_{z_j}(\omega) \stackrel{\Delta}{=} \mathcal{F}[\mathbf{z}_j(t)] \quad (7)$$

$$\mathbf{S}_{u_i}(\omega) \stackrel{\Delta}{=} \mathcal{F}[\mathbf{u}_i(t)]$$

The asterisk $*$ denotes complex conjugation, ω is frequency in radians, and \mathcal{F} is a discrete Fourier transform (this work used a fast Fourier transform or FFT). The input force in the time domain at location j is $\mathbf{z}_j(t)$, and the output response (i.e., acceleration, velocity, or displacement) in the time domain at i is $\mathbf{u}_i(t)$.

With assumed linear behaviour and proportional damping, the accelerance FRF (i.e., given acceleration response data) can also be estimated using modal parameters,

$$\mathbf{H}_{ij}(\omega) = -\omega^2 \sum_{k=1}^n \frac{\mathbf{A}_{ij}^{(k)}}{\omega_{nk}^2 - \omega^2 + 2i\zeta_k \omega \omega_{nk}} \quad (8)$$

where $\mathbf{A}_{ij}^{(k)}$ is the residue for mode k , defined as the product of the mass-normalised mode shapes at locations i and j ($\mathbf{A}_{ij}^{(k)} = \psi_{ik} \psi_{jk}$) [17]. The natural frequency associated with mode k is ω_{nk} , and the modal damping associated with mode k is ζ_k [17]. In the examples below, Eq. (8) was used to compute a non-zero prior mean function for the (OM)GP form.

2.2. Conventional Gaussian process regression

Gaussian processes provide a Bayesian nonparametric approach to regression/classification and are a suitable basis for a population form because of their robust uncertainty quantification. For a set of training data $D = \{x_i, y_i\}_{i=1}^N = \{\mathbf{x}, \mathbf{y}\}$, where x_i and y_i are sets of N inputs and outputs, respectively, the GP regression attempts to obtain the predictive distribution for observation y_* given a new input x_* and training data D [13]. The GP regression model can be expressed as a noiseless latent function $f(x_i)$ plus an independent noise term ϵ_i [13], written as,

$$y_i = f(x_i) + \epsilon_i \quad (9)$$

A GP prior is placed over the latent function $f(x_i)$ and a Gaussian prior is placed over the noise term ϵ_i ,

$$f(x_i) \sim \mathcal{GP}(m(x_i), k(x_i, x_j)), \epsilon_i \sim \mathcal{N}(0, \sigma^2) \quad (10)$$

where $m(x_i)$ is a mean function, $k(x_i, x_j)$ is a covariance function, and σ^2 is a noise variance term and the first hyperparameter. Over a finite, arbitrary set of inputs $\mathbf{x} = \{x_1, \dots, x_N\}$, the GP is a (joint) multivariate Gaussian [11,18],

$$p(\mathbf{f} | \mathbf{x}) = \mathcal{N}(\mathbf{m}, \mathbf{K}_{\mathbf{xx}}) \quad (11)$$

where $\mathbf{m} = \{m(x_1), \dots, m(x_N)\}$, $\mathbf{f} = \{f(x_1), \dots, f(x_N)\}$, and $\mathbf{K}_{\mathbf{xx}}$ is the covariance matrix, such that $\mathbf{K}_{\mathbf{xx}}[i, j] = k(x_i, x_j) \forall i, j \in \{1, \dots, N\}$ [11]. Note that in certain instances within this paper, square brackets are used to index matrices and vectors for readability.

The mean function $m(x_i)$ is often assumed to be zero, given that the GP is flexible enough to model arbitrary trends [11,18,19]. However, the GP fit may be improved, by incorporating *a priori* knowledge from engineering judgement or an understanding of the system physics, into the GP prior by choosing suitable mean and covariance functions [3,11]. Coupling behaviour between y_i and y_j , and therefore important properties such as process variance and smoothness, is specified by the covariance function [13]. Appropriate covariance functions are selected according to the shape and smoothness of the data. This work applied the commonly-used squared exponential,

$$k(x_i, x_j) = \sigma_f^2 \exp\left(-\frac{1}{2l^2}(x_i - x_j)^2\right) \quad (12)$$

where σ_f^2 and l are second and third hyperparameters, representing the process variance and length scale, respectively. The length scale l influences how fast the correlation between outputs decays according to the input separation, and the process variance σ_f^2 determines the variance of the modelled signal [13].

The remaining hyperparameters α are those required by the mean function, if non-zero, giving the full model hyperparameters as $\theta = \{l, \sigma_f, \sigma, \alpha\}$. In this work, the modal damping-based FRF estimation from Eq. (8) provided the mean function for the GP regression. Modal damping (ζ_k), residue ($A_{ij}^{(k)}$), and natural frequency (ω_{nk}), provided the mean-function hyperparameters. Keeping θ fixed, the joint distribution of the training data $D = \{\mathbf{x}, \mathbf{y}\} = \{x_i, y_i\}_{i=1}^N$ and unknown function values at test locations under the prior $\{\mathbf{x}_*, \mathbf{y}_*\} = \{\mathbf{x}_*[i], \mathbf{y}_*[i]\}_{i=1}^M$ is a multivariate Gaussian, written as,

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{y}_* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{m} \\ \mathbf{m}_* \end{bmatrix}, \begin{bmatrix} \mathbf{K}_{\mathbf{xx}} + \mathbf{R} & \mathbf{K}_{\mathbf{xx}_*} \\ \mathbf{K}_{\mathbf{x}_*\mathbf{x}} & \mathbf{K}_{\mathbf{x}_*\mathbf{x}_*} + \mathbf{R}_* \end{bmatrix}\right) \quad (13)$$

$$\begin{aligned} \mathbf{R} &\stackrel{\Delta}{=} \sigma^2 \mathbf{I}_N \\ \mathbf{R}_* &\stackrel{\Delta}{=} \sigma^2 \mathbf{I}_M \end{aligned} \quad (14)$$

where \mathbf{I} is an identity matrix of dimension N or M as specified [11,19]. Likewise, $\mathbf{m}_* = \{m(\mathbf{x}_*[i])\}_{i=1}^M$ is the mean vector for new observations [11].

Conditioning the joint distribution in Eq. (13) on the training data in D yields the predictive distribution over \mathbf{y}_* , with mean $\boldsymbol{\mu}_*$ and covariance $\boldsymbol{\Sigma}_*$ [19], defined as,

$$\begin{aligned} p(\mathbf{y}_* | \mathbf{x}_*, D) &= \mathcal{N}(\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*) \\ \boldsymbol{\mu}_* &\stackrel{\Delta}{=} \mathbf{m}_* + \mathbf{K}_{\mathbf{x}_*\mathbf{x}} (\mathbf{K}_{\mathbf{xx}} + \mathbf{R})^{-1} (\mathbf{y} - \mathbf{m}) \\ \boldsymbol{\Sigma}_* &\stackrel{\Delta}{=} \mathbf{K}_{\mathbf{x}_*\mathbf{x}_*} - \mathbf{K}_{\mathbf{x}_*\mathbf{x}} (\mathbf{K}_{\mathbf{xx}} + \mathbf{R})^{-1} \mathbf{K}_{\mathbf{xx}} + \mathbf{R}_* \end{aligned} \quad (15)$$

This work applied a Type-II maximum likelihood approach to learn the hyperparameters, which maximises the marginal likelihood of the model $p(\mathbf{y} | \mathbf{x}, \theta)$ [19]. This procedure involves minimising the negative log-marginal-likelihood to improve numerical stability [18]. Therefore, the hyperparameters were optimised via,

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \{-\log p(\mathbf{y} | \mathbf{x}, \theta)\} \quad (16)$$

where,

$$\begin{aligned} -\log p(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}) &= -\log \mathcal{N}(\mathbf{y} | \mathbf{m}, \mathbf{K}_{\mathbf{x}\mathbf{x}} + \mathbf{R}) \\ &= \frac{N}{2} \log(2\pi) + \frac{1}{2} \log |\mathbf{K}_{\mathbf{x}\mathbf{x}} + \mathbf{R}| \\ &\quad + \frac{1}{2} [(\mathbf{y} - \mathbf{m})^T (\mathbf{K}_{\mathbf{x}\mathbf{x}} + \mathbf{R})^{-1} (\mathbf{y} - \mathbf{m})] \end{aligned} \quad (17)$$

Constrained nonlinear optimisation was performed via the *fmincon* function in MATLAB, and appropriate bounds were applied to the mean-function hyperparameters, given prior knowledge of the modal characteristics of the helicopter blades.

2.3. An overlapping mixture of Gaussian processes (OMGP)

In one example, this work used the conventional GP equations to evaluate a supervised mixture of GPs by assuming known which data belonged to each helicopter blade, and fitting each data set separately. In a second example, the data were fitted using an overlapping mixture of Gaussian processes (OMGP), where the number of trajectory functions were assumed known *a priori*, but it was not assumed which data belonged to a given class. In this case, the unknown class was the particular helicopter blade that produced the measurements. In practice, it is likely that this information would be known; however, there are regularly other forms of missing data labels, such as changes in operating conditions, or possible damage. The helicopter blades case study is also representative of these types of situations, commonly encountered in PBSHM, where certain information is unavailable. As it is an unsupervised learning algorithm, the OMGP is a useful and flexible method for developing the population form in the absence of complete information.

The OMGP [3,11,13,20], describes a data set by evaluating each observation using one of K latent functions plus additive noise,

$$y_i^{(k)} = \{f^{(k)}(x_i) + \epsilon_i\}_{k=1}^K \quad (18)$$

The OMGP is unsupervised, i.e., the labels that assign the data to a given function are unknown. Therefore, another latent variable is introduced, \mathbf{Z} , which is a binary indicator matrix. The matrix \mathbf{Z} , with entries $\mathbf{Z}[i, k]$, is entirely populated with zeros, except for one non-zero entry per row. The non-zero entries, $\mathbf{Z}[i, k] \neq 0$, indicate that observation i was generated by function k . The likelihood of the OMGP is therefore written as [11,13],

$$p(\mathbf{y} | \{\mathbf{f}^{(k)}\}_{k=1}^K, \mathbf{Z}, \mathbf{x}) = \prod_{i,k=1}^{N,K} p(y_i | f^{(k)}(x_i))^{\mathbf{Z}[i,k]} \quad (19)$$

Prior distributions are then placed over the latent functions and variables, such that [11],

$$P(\mathbf{Z}) = \prod_{i,k=1}^{N,K} \mathbf{\Pi}[i, k]^{\mathbf{Z}[i,k]} \quad (20)$$

$$f^{(k)}(x_i) \sim \mathcal{GP}(m^{(k)}(x_i), k^{(k)}(x_i, x_j)), \epsilon_i \sim \mathcal{N}(0, \sigma^2) \quad (21)$$

where $P(\mathbf{Z})$ is the prior over the indicator matrix \mathbf{Z} , and $\mathbf{\Pi}[i, :]$ is a histogram over the K components for the i th observation, and $\sum_{k=1}^K \mathbf{\Pi}[i, k] = 1$ [11]. The terms in Eq. (21) are independent GP priors over each latent function $f^{(k)}$ with distinct mean $m^{(k)}(x_i)$ and kernel $k^{(k)}(x_i, x_j)$ functions [11]. A shared hyperparameter σ , is used to define the prior over the noise variances, to reduce the number of latent variables in the calculations [11]. As with the conventional GP, this work used the modal damping-based FRF estimation from Eq. (8) as the mean function for the GP regression. Modal damping, residue, and natural frequency provided the mean-function hyperparameters.

2.3.1. Variational inference

Exact computation of the posterior distribution $p(\{\mathbf{f}^{(k)}\}_{k=1}^K, \mathbf{Z} | \mathbf{x}, \mathbf{y})$ is intractable; therefore, a variational inference [21] and Expectation Maximisation (EM) scheme was employed to estimate the posterior, by specifying an approximate density family $q(\mathbf{a}) \in \mathcal{Q}$ over the target conditional $p(\mathbf{a} | \mathbf{b})$ [3,11,13,21]. The best approximation $\hat{q}(\mathbf{a})$ of the true posterior distribution $p(\mathbf{a} | \mathbf{b})$ can be expressed via the KL-divergence [11,21],

$$\hat{q}(\mathbf{a}) = \underset{q(\mathbf{a}) \in \mathcal{Q}}{\operatorname{argmin}} \{\operatorname{KL}(q(\mathbf{a}) \parallel p(\mathbf{a} | \mathbf{b}))\} \quad (22)$$

$$\begin{aligned} \operatorname{KL}(q(\mathbf{a}) \parallel p(\mathbf{a} | \mathbf{b})) &= \mathbb{E}_{q(\mathbf{a})}[\log q(\mathbf{a})] - \mathbb{E}_{q(\mathbf{a})}[\log p(\mathbf{a} | \mathbf{b})] \\ &= \mathbb{E}_{q(\mathbf{a})}[\log q(\mathbf{a})] - \mathbb{E}_{q(\mathbf{a})}[\log p(\mathbf{a}, \mathbf{b})] \\ &\quad + \log p(\mathbf{b}) \end{aligned} \quad (23)$$

where, for this work, $\mathbf{a} \triangleq \{\{\mathbf{f}^{(k)}\}, \mathbf{Z}\}$ and $\mathbf{b} \triangleq \{\mathbf{y}\}$. However, $p(\mathbf{b})$ is intractable [11,21]; therefore, rather than the KL-divergence, an alternative object is defined, called the *evidence lower bound* or *elbo*. The elbo, \mathcal{L}_b , is equivalent to the negative KL-divergence in Eq. (23), up to the term $\log p(\mathbf{b})$, which is a constant with respect to $q(\mathbf{z})$ [11],

$$\begin{aligned} \mathcal{L}_b(\mathbf{a}) &= \mathbb{E}_{q(\mathbf{a})}[\log p(\mathbf{a}, \mathbf{b})] - \mathbb{E}_{q(\mathbf{a})}[\log q(\mathbf{a})] \\ &= \int q(\mathbf{a}) \log \frac{p(\mathbf{a}, \mathbf{b})}{q(\mathbf{a})} \mathbf{d}\mathbf{a} \end{aligned} \quad (24)$$

Maximising \mathcal{L}_b minimises the KL-divergence between $q(\mathbf{a})$ and $p(\mathbf{a} | \mathbf{b})$. Rearranging Eq. (23) and substituting in Eq. (24) yields,

$$\log p(\mathbf{b}) = \text{KL}(q(\mathbf{a}) \parallel p(\mathbf{a} | \mathbf{b})) + \mathcal{L}_b \quad (25)$$

which shows that because $\text{KL}(\cdot) \geq 0$ [22], the evidence is lower-bounded by the elbo \mathcal{L}_b , i.e., $\log p(\mathbf{b}) \geq \mathcal{L}_b$ [11]. Substituting $\mathbf{a} \stackrel{\Delta}{=} \{\{\mathbf{f}^{(k)}\}, \mathbf{Z}\}$ and $\mathbf{b} \stackrel{\Delta}{=} \{\mathbf{y}\}$ into Eq. (24) yields,

$$\begin{aligned} \log p(\mathbf{y} | \mathbf{x}) &= \log \iint p(\{\mathbf{f}^{(k)}\}, \mathbf{Z}, \mathbf{y}, \mathbf{x}) p(\{\mathbf{f}^{(k)}\}), p(\mathbf{Z}) d\{\mathbf{f}^{(k)}\} d\mathbf{Z} \\ &\geq \mathcal{L}_b = \iint q(\{\mathbf{f}^{(k)}\}, \mathbf{Z}) \log \frac{p(\{\mathbf{f}^{(k)}\}, \mathbf{Z}, \mathbf{y}, \mathbf{x})}{q(\{\mathbf{f}^{(k)}\}, \mathbf{Z})} d\{\mathbf{f}^{(k)}\} d\mathbf{Z} \\ &= \iint q(\{\mathbf{f}^{(k)}\}, \mathbf{Z}) \log \frac{p(\mathbf{y} | \{\mathbf{f}^{(k)}\}, \mathbf{Z}, \mathbf{x}) p(\mathbf{Z}) \prod_{k=1}^K p(\mathbf{f}^{(k)} | \mathbf{x})}{q(\{\mathbf{f}^{(k)}\}, \mathbf{Z})} d\{\mathbf{f}^{(k)}\} d\mathbf{Z} \end{aligned} \quad (26)$$

An improved bound, \mathcal{L}_{bc} , that lower-bounds the marginal likelihood while also being an upper-bound on the standard variational bound \mathcal{L}_b [11,23], was implemented in this work. The improved lower bound can be expressed as [11,24],

$$\begin{aligned} \log p(\mathbf{y} | \mathbf{x}) &\geq \mathcal{L}_{bc} \\ &= \sum_{k=1}^K \left(\frac{1}{2} \left\| \mathbf{R}^{(k)\top} \setminus \left(\mathbf{B}^{(k)\frac{1}{2}} (\mathbf{y} - \mathbf{m}^{(k)}) \right) \right\|^2 - \sum_{i=1}^N \log \mathbf{R}^{(k)}[i, i] \right) \dots \\ &\quad - \text{KL} \left(q(\mathbf{Z}) \parallel p(\mathbf{Z}) \right) - \frac{1}{2} \sum_{i=1, k=1}^{N, K} \hat{\mathbf{H}}[i, k] \log (2\pi\sigma^2) \\ \mathbf{R}^{(k)} &\stackrel{\Delta}{=} \text{chol} \left(\mathbf{I} + \mathbf{B}^{(k)\frac{1}{2}} \mathbf{K}_{xx}^{(k)} \mathbf{B}^{(k)\frac{1}{2}} \right) \\ \text{KL} \left(q(\mathbf{Z}) \parallel p(\mathbf{Z}) \right) &\stackrel{\Delta}{=} \sum_{i=1, k=1}^{N, K} \hat{\mathbf{H}}[i, k] \log \frac{\hat{\mathbf{H}}[i, k]}{\hat{\mathbf{H}}[i, k]} \end{aligned} \quad (27)$$

where the backslash operator solves the system of linear equations $\mathbf{A}\mathbf{x} = \mathbf{B}$ for \mathbf{x} via $\mathbf{A} \setminus \mathbf{B}$ and $\text{chol}(\cdot)$ is the Cholesky decomposition. The quantity $\mathbf{B}^{(k)}$ is an $N \times N$ diagonal matrix with elements,

$$\mathbf{B}^{(k)} = \text{diag} \left(\left\{ \frac{\hat{\mathbf{H}}[1, k]}{\sigma^2}, \dots, \frac{\hat{\mathbf{H}}[N, k]}{\sigma^2} \right\} \right) \quad (28)$$

The improved lower bound \mathcal{L}_{bc} depends only on $p(\mathbf{Z})$ and is independent of $p(\{\mathbf{f}^{(k)}\})$; as such, \mathcal{L}_{bc} can be referred to as the *marginalised variational bound* [11,13]. For a given $p(\mathbf{Z})$, when an optimal choice for $p(\{\mathbf{f}^{(k)}\})$ is made, \mathcal{L}_{bc} is equivalent to \mathcal{L}_b [11,13]. This idea means that the bound is more stable over a variety of hyperparameters, because the improved bound is independent of $p(\{\mathbf{f}^{(k)}\})$ [11,13,23]. Using the improved lower bound is therefore useful in the M-step of the Expectation Maximisation (EM) scheme, when \mathcal{L}_{bc} is optimised with respect to the hyperparameters until convergence, with $p(\mathbf{Z})$ held fixed [13].

2.3.2. E-step: Mean-field updates

A mean-field assumption was employed, and a family of variational distributions $q(\mathbf{a}) \in \mathcal{Q}$ were defined such that each latent variable could be optimised independently, i.e., $q(\{\mathbf{f}^{(k)}\}, \mathbf{Z}) = q(\{\mathbf{f}^{(k)}\})q(\mathbf{Z})$. This assumption allowed for the optimisation of \mathcal{L}_b (or \mathcal{L}_{bc}) with respect to a given latent variable while keeping the others fixed, and iteratively updating each variable until convergence of the lower bound was reached (i.e., EM Maximisation) [11].

If $q(\{\mathbf{f}^{(k)}\})$ and the marginals for each component $q(\mathbf{f}^{(k)}) = \mathcal{N}(\boldsymbol{\mu}^{(k)}, \boldsymbol{\Sigma}^{(k)})$ are assumed known, \mathcal{L}_b can be analytically maximised with respect to $q(\mathbf{Z})$ by constraining q to be a probability density and setting the derivative of the bound to zero [11,13],

$$\begin{aligned} q(\mathbf{Z}) &= \prod_{i, k=1}^{N, K} \hat{\mathbf{H}}[i, k] \mathbf{Z}^{[i, k]}, \text{ s.t. } \hat{\mathbf{H}}[i, k] \propto \mathbf{H}[i, k] \exp(a_{ik}) \\ a_{ik} &\stackrel{\Delta}{=} -\frac{1}{2\sigma^2} \left(\left(y_i - \boldsymbol{\mu}_i^{(k)} \right)^2 + \boldsymbol{\Sigma}^{(k)}[i, i] \right) - \frac{1}{2} \log (2\pi\sigma^2) \end{aligned} \quad (29)$$

where Eq. (29) implies that $q(\mathbf{Z})$ is factorised for each sample [11,13].

Now, assuming that $q(\mathbf{Z})$ is known, the lower bound can be analytically maximised with respect to $q(\{\mathbf{f}^{(k)}\})$, as in,

$$\begin{aligned} q(\mathbf{f}^{(k)}) &= \mathcal{N} \left(\mathbf{f}^{(k)} \mid \boldsymbol{\mu}^{(k)}, \boldsymbol{\Sigma}^{(k)} \right) \\ \boldsymbol{\Sigma}^{(k)} &\stackrel{\Delta}{=} \left(\mathbf{K}_{xx}^{-1(k)} + \mathbf{B}^{(k)} \right)^{-1} \\ \boldsymbol{\mu}^{(k)} &\stackrel{\Delta}{=} \mathbf{m}^{(k)} + \boldsymbol{\Sigma}^{(k)} \mathbf{B}^{(k)} (\mathbf{y} - \mathbf{m}^{(k)}) \end{aligned} \quad (30)$$

The best candidate $\hat{q}(\mathbf{a})$ that most closely approximates the true posterior distribution was found by initialising $q(\mathbf{Z})$ and $q(\{\mathbf{f}^{(k)}\})$ from their respective priors, and iteratively updating them by alternating Eqs. (29) and (30). As both updates are optimal with respect to the distribution that they compute, they are guaranteed to increase the lower bound on the log-marginal-likelihood [11,13], computed via Eq. (26).

2.3.3. M-step: Hyperparameter optimisation

In the E-step, \mathcal{L}_{bc} was optimised with respect to $p(\mathbf{Z})$ until convergence by iterating Eqs. (29) and (30), with the hyperparameters held fixed. In the M-step, the lower bound is optimised with respect to all model hyperparameters, as in,

$$\left\{ \{\hat{\theta}_k\}_{k=1}^K, \hat{\boldsymbol{\Pi}} \right\} = \underset{\{\{\theta_k\}_{k=1}^K, \boldsymbol{\Pi}\}}{\operatorname{argmax}} \{ \mathcal{L}_{bc} \} \quad (31)$$

while keeping the approximate posterior $q(\mathbf{Z})$ fixed. The E- and M-steps are then alternated until the lower bound converges. In this work, constrained nonlinear optimisation was performed via the *fmincon* function in MATLAB, and appropriate bounds were applied to the mean-function hyperparameters given prior knowledge of the modal characteristics of the helicopter blades.

2.3.4. Predictive equations

Once learnt, the OMGP fit (or form) can be used to estimate the latent variables and functions and to inform damage detection. For training data D , the *maximum a posteriori* (MAP) estimate can be used to categorise observations according to the most likely component k , via [11],

$$\hat{k}_i = \underset{k}{\operatorname{argmax}} \{ \hat{\boldsymbol{\Pi}}[i, k] \} \quad (32)$$

The MAP class component for new data \hat{k}_* is then defined as [11],

$$\hat{k}_* = \underset{k_*}{\operatorname{argmax}} \{ p(k_* | \mathbf{x}_*, \mathbf{y}_*, D) \} \quad (33)$$

where, for a set of test data, the posterior predictive class component k_* is expressed using Bayes Rule (note that to classify new data according to k_* , both \mathbf{x}_* and \mathbf{y}_* must be observed [11]),

$$p(k_* | \mathbf{x}_*, \mathbf{y}_*, D) = \frac{p(\mathbf{y}_* | \mathbf{x}_*, k_*, D)p(k_*)}{p(\mathbf{y}_* | \mathbf{x}_*, D)} \quad (34)$$

The numerator of Eq. (34), or unnormalised posterior, is defined as [11],

$$p(\mathbf{y}_* | \mathbf{x}_*, k, D)p(k_*) \stackrel{\Delta}{=} \mathcal{N}(\mathbf{y}_* | \boldsymbol{\mu}_*^{(k)}, \boldsymbol{\Sigma}_*^{(k)}) \boldsymbol{\Pi}[*], k] \quad (35)$$

and the denominator of Eq. (34), which is the marginal likelihood (evidence), is [11],

$$\begin{aligned} p(\mathbf{y}_* | \mathbf{x}_*, D) &\approx \sum_{k=1}^K \boldsymbol{\Pi}[*], k] \int p(\mathbf{y}_* | \mathbf{f}^{(k)}, \mathbf{x}_*, D) q(\mathbf{f}^{(k)} | D) d\mathbf{f}^{(k)} \\ &= \sum_{k=1}^K \mathcal{N}(\mathbf{y}_* | \boldsymbol{\mu}_*^{(k)}, \boldsymbol{\Sigma}_*^{(k)}) \boldsymbol{\Pi}[*], k] \end{aligned} \quad (36)$$

where [11],

$$\begin{aligned} \boldsymbol{\mu}_*^{(k)} &\stackrel{\Delta}{=} \mathbf{m}_*^{(k)} + \mathbf{K}_{\mathbf{x}_* \mathbf{x}}^{(k)} \left(\mathbf{K}_{\mathbf{x}\mathbf{x}}^{(k)} + \mathbf{B}^{(k)-1} \right)^{-1} (\mathbf{y} - \mathbf{m}^{(k)}) \\ \boldsymbol{\Sigma}_*^{(k)} &\stackrel{\Delta}{=} \mathbf{K}_{\mathbf{x}_* \mathbf{x}_*}^{(k)} - \mathbf{K}_{\mathbf{x}_* \mathbf{x}}^{(k)} \left(\mathbf{K}_{\mathbf{x}\mathbf{x}}^{(k)} + \mathbf{B}^{(k)-1} \right)^{-1} \mathbf{K}_{\mathbf{x}\mathbf{x}_*}^{(k)} + \mathbf{R}_*^{(k)} \\ \mathbf{R}_*^{(k)} &\stackrel{\Delta}{=} \sigma^2 \mathbf{I}_M \end{aligned} \quad (37)$$

The prior mixing proportions for new observations, $\boldsymbol{\Pi}[*], k]$, weight each component equally *a priori*, such that the sum of the weights is equal to 1, or $\boldsymbol{\Pi}[*], k] = 1/K$ [11]. The predictive equations for the OMGP are quite similar to those for the conventional GP in Eq. (15), with the exception of the noise component for the training data $\mathbf{B}^{(k)-1}$ which is scaled for each sample according to the inverse of the mixing proportion $\hat{\boldsymbol{\Pi}}[i, k]^{-1}$ [11,13]. As such, the contribution of each sample is weighted with respect to its posterior predictive component in the mixture [11].

2.3.5. Practical implementation of the described technology

In this work, a population form was developed for the helicopter blades using a mixture of Gaussian processes, and new (simulated) FRF data were used to evaluate the sensitivity of the form for novelty detection. There were four healthy helicopter blades and therefore four classes ($K = 4$). Because the new data had an equal chance of belonging to any of the four classes, the prior mixing proportion weighted each class equally, with $\boldsymbol{\Pi}[*], k] = 0.25$. The marginal likelihood (evidence), defined in Eq. (36) was used as a novelty index, and provided an averaged likelihood of new data belonging to any of the k classes. For computational ease and to avoid numerical underflow, the calculations were performed in log space. Specifically, negative log quantities were used, such that the visual presentation of data novelty was consistent with the Mahalanobis squared-distance (MSD) novelty index from [3,4,12]. As such, data below a certain threshold were considered normal or *inlying*, and data above the threshold were considered novel or *outlying*.

A suitable threshold for novelty detection was determined using a similar method as with the MSD in [3,4,12]. Bootstrap-sampling was used, where 1000 samples were randomly selected from the normal-condition data used to train the form. The negative log-marginal-likelihood was then calculated for the samples, for a large number of trials. The critical value was the threshold with a certain percentage of the calculated values below it. This work used a 99% confidence interval, or 2.58σ .

2.3.6. Model evaluation

The normalised mean squared-error (NMSE) and Mahalanobis squared-distance (MSD) were calculated to compare the different models developed in this work (e.g., supervised versus unsupervised mixture of GPs), particularly when it was not feasible to use the model's objective function as a performance metric (e.g., when comparing two models learnt using different training data). The NMSE was computed to assess each test observation against the mean of its most likely class component k_* , via [11],

$$\text{NMSE} = \frac{100}{M\sigma_{y_*}^2} \left(\boldsymbol{\mu}_*^{(k_*)} - \mathbf{y}_* \right)^T \left(\boldsymbol{\mu}_*^{(k_*)} - \mathbf{y}_* \right) \quad (38)$$

Likewise, the (normalised) MSD was computed via,

$$\text{MSD} = \frac{1}{M} \sum_{i=1}^M \frac{\left(\boldsymbol{\mu}_*^{(k_*)}[i] - y_*[i] \right)^2}{\boldsymbol{\Sigma}_*^{(k_*)}[i, i]} \quad (39)$$

where M is the number of test data points (not included in training set). In this work, the MSD was used only to assess the predictive variance $\boldsymbol{\Sigma}_*^{(k)}$. The MSD is less useful when assessing the fit of the OMGP, as the error is scaled by the predictive variance $\boldsymbol{\Sigma}_*^{(k)}$ [11].

3. Experimental campaign on a population of helicopter blades

Vibration data were collected over a series of tests from a set of four healthy, full-scale composite helicopter blades, in fixed–free and free–free configurations. This section discusses how differences among the blades and changes in boundary conditions present as changes in the FRFs of the blades, particularly with respect to the locations of the peaks.

3.1. Experiments


Vibration data were collected at ambient temperature from four healthy, full-scale composite helicopter blades using Siemens PLM LMS SCADAS hardware and software at the Laboratory for Verification and Validation (LVV) in Sheffield, UK. Testing was performed in fixed–free and free–free configurations. To approximate a fixed–free boundary condition, the root end of each blade was placed in a substantiated strong-wall mount. To approximate a free–free boundary condition, the blade was suspended from a gantry via springs and cables.

Ten uniaxial 100 mV/g accelerometers were placed along the length of the underside of each blade. Note that the same accelerometers were used on each blade, and care was taken to ensure that they were attached to approximately the same locations on each blade. For the fixed–free tests, an electrodynamic shaker with force gauge was mounted to a fixture bolted to the laboratory floor and attached to the blade 0.575 metres from the fixed end. The shaker was attached to the underside of the blade in the flapwise direction. A continuous random excitation was generated in LMS (note: LMS refers to Siemens PLM LMS SCADAS hardware and software) and applied to excite the blade up to 400 Hz. Approximately 7.4 min of throughput force and acceleration data were collected for each test, with a time step of 1.25e–03 s and a sample rate of 800 Hz. The data were then divided into 20 blocks, each with a dimension of 16,384. Some data at the start and end of the acquisition were discarded as these data corresponded to powering up and down of the shaker. A Hanning window was applied to each data block and FRFs were computed in LMS. The FRFs were then averaged in the frequency domain. For the free–free tests, the shaker was connected at the same location on the blade as the fixed–free tests, but was suspended from a gantry via springs. A continuous random excitation was generated in LMS and applied to excite the blade up to 512 Hz. Approximately 5.8 min of throughput force and acceleration data were collected for each test, with a time step of 9.77e–04 s and a sample rate of 1024 Hz. As with the fixed–free tests, the data were divided into 20 blocks, each with a dimension of 16,384, and some data corresponding to powering up/powering down of the shaker were discarded. A Hanning window was applied to each data block and FRFs were computed for each measurement and then averaged in the frequency domain. The measured bandwidths for the fixed–free and free–free tests were selected to maximise the number and accuracy of the captured modes (there was a large dip in the input spectrum just after 400 Hz for the fixed–free tests and just after 512 Hz for the free–free tests). Data were collected over a relatively-long time period for each test and a large number of averages were used to improve the spectra at the lower modes near the limits of the operating range of the sensors. The experimental setup for the fixed–free and free–free tests are shown in Figs. 1(a) and 1(b), respectively. The sensor layout is shown in Fig. 2. Acquisition parameters are listed in Table 1.

An example drive-point coherence spectrum and FRF for Blade 3 in fixed–free and free–free conditions (all computed from data obtained at the accelerometer located near the force-input location, 0.575 metres from the blade root) are shown in Figs. 3(a) and 3(b), respectively. These spectra are representative of all measurements and show that coherence values were close to 1 at the resonances, with the exception of some difficulty with collecting the first two modes below 10 Hz. This difficulty in acquiring the first few modes resulted in part from the operating frequency range of the accelerometers and the extremely high flexibility of the blades, which limited the shaker placement.



(a)



(b)

Fig. 1. Helicopter blade in (a) fixed-free and (b) free-free configurations.




Fig. 2. Sensor locations on helicopter blades.

Table 1
Acquisition parameters for vibration testing of helicopter blades.

Acquisition settings	Fixed-free	Free-free
Time step	1.25e-03 s	9.77e-04 s
Acquisition time	20.48 s	16.00 s
Bandwidth	400 Hz	512 Hz
Lines	8192	8192
Frequency step	4.88e-02 Hz	6.25e-02 Hz
Window	Hanning	Hanning

3.2. Discussion of experimental results

This work is concerned with demonstrating how differences among the nominally-identical blades (with consideration for boundary condition effects), present as changes in the peak positions as visible in the FRFs. Data were collected from 10 accelerometers during each test. This discussion considers the results from the accelerometer near the blade tip, as the feature changes visible in these spectra are representative of all test points along the blade length. The averaged FRFs for each blade in the fixed-free condition are shown in Figs. 4(a) and 4(b). Fig. 4(a) shows the full measured 400 Hz bandwidth and Fig. 4(b) shows only modes below 80 Hz. Likewise, the averaged FRF for each blade in free-free is shown in Figs. 5(a) and 5(b), where Fig. 5(a) shows the full measured 512 Hz bandwidth and Fig. 5(b) shows only modes below 80 Hz.




Fig. 3. Representative drive-point coherence and FRF from Blade 3 in (a) fixed-free and (b) free-free configurations.

Figs. 4(a) and 4(b) show increasing variability with respect to frequency, which is an expected result, given that higher-frequency modes are more sensitive to small physical changes than lower-frequency modes. For modes less than 80 Hz, the maximum frequency difference among the blades was approximately 2.5 Hz, while for modes greater than 80 Hz, the maximum frequency difference was approximately 6.3 Hz. Figs. 5(a) and 5(b) show a very different distribution of peaks compared to the fixed-free tests (as expected, given such a dramatic change in boundary condition), but also show similar variation among the blades. For modes less than 80 Hz, the maximum frequency difference among the blades was approximately 1.6 Hz. For modes greater than 80 Hz, the maximum frequency difference among the blades was approximately 6.8 Hz.

Note that both the fixed-free and free-free results show grouping at several of the peaks, where Blades 1 and 2 appear closely aligned in frequency while Blades 3 and 4 appear closely aligned. These results are particularly noticeable in Fig. 4(b), between 40 Hz and 60 Hz, and in Fig. 5(b), between 60 Hz and 80 Hz. These findings are quite relevant for PBSHM, which seeks to transfer valuable information across similar structures [6]. If these data were used to develop a normal condition for the blades against which to evaluate new data for novelty, this variation could present some interesting problems (e.g., negative transfer). For example, if Blade 4 experienced some damage that reduced its natural frequency, such that it aligned with the same peak from Blade 1, the damage may be missed. Another potential issue would be if a new blade were added to the population and was slightly less stiff than Blade 2, the new blade may be flagged as damaged when compared to the normal condition developed using the original blades and test results. The development of a population form to represent the behaviour of the group is discussed in the next section. Only the fixed-free results were used in developing the form, although significant changes in boundary conditions could be addressed using techniques such as transfer learning [6]. Both the fixed-free and free-free tests showed considerable variability among the blades and a similar grouping pattern.

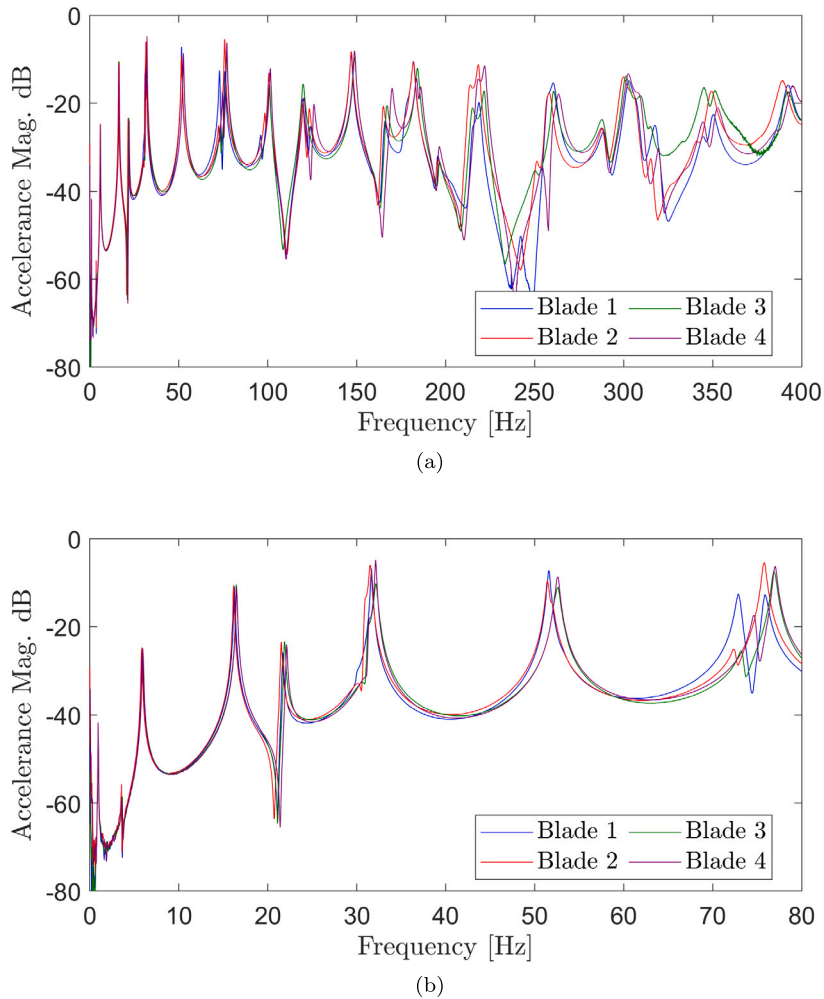


Fig. 4. FRF magnitudes for the helicopter blades in fixed-free (a) full bandwidth and (b) first 80 Hz.

4. Development of a Gaussian process form

To differentiate between normal and damaged states, an important development in PBSHM will involve establishing a normal condition for a population of structures. In certain cases, the behaviour of the group can be represented using a general model, called a *form*, against which new data can be evaluated for novelty [3,4,11]. This work used Gaussian process (GP) regression to model the form for the FRFs of the helicopter blades in fixed-free (recall that the free-free tests showed similar variability among the blades as the fixed-free tests).

Two examples are discussed. In the first example, the form was developed using a supervised mixture of GPs with normal-condition experimental data from the blades. This example was supervised because it was assumed *a priori* which data belonged to each helicopter blade. In the second example, the normal-condition form was developed using an (unsupervised) OMGP, where the number of trajectory functions were assumed known *a priori*, but it was not assumed which data belonged to a given helicopter blade, representing a situation with incomplete label information. This situation often occurs in SHM, when the labels correspond to a range of operational, environmental, or damage conditions for each structure.

4.1. A form using a supervised mixture of Gaussian processes

A supervised mixture of Gaussian processes assumes that the data labels are known *a priori*. In this case, it was assumed to be known which FRF data were associated with each helicopter blade. This assumption allowed for each GP in the mixture to be treated independently using the equations for a conventional GP (Eqs. (9)–(17)). While in a testing environment, FRFs are often described via magnitude and phase; however, the uncertainty distribution of the magnitude is non-Gaussian. As such, it is not proper to fit the magnitude directly using GPs (doing so will result in support for negative values, but $\text{mag}[H_{ij}] \geq 0$). If the time domain data are




Fig. 5. FRF magnitudes for the helicopter blades in free-free (a) full bandwidth and (b) first 80 Hz.

random realisations of a stochastic process, the fast Fourier transforms (FFT) of the signals are random realisations of the spectra, because the FFT is a linear operator [25]. As quotients of FFTs (rather, quotients of cross-spectra and input auto-spectra), the real and imaginary parts of the FRF are not strictly Gaussian; however, for this paper it is assumed that each can be approximated as GPs. In this work, the FRF was expressed as real and imaginary parts, and the real and imaginary parts of the FRF were fit separately via Eqs. (9)–(17). Note that several conference papers related to this research [26,27], provided a preliminary approach where magnitude was fitted directly. Although the resulting form was useful for damage detection, the support for negative magnitudes does not make sense and was therefore not employed in later work.

To develop the form, a narrow frequency band was selected between 48 and 56 Hz, containing the fifth bending mode of the blades. Closer inspection of the data revealed a likely second mode in the same frequency range; however, the modes were so closely-spaced that they could conceivably be a slightly distorted single mode. An SDOF assumption was imposed, and an SDOF mean function was applied to the GPs. In other words, the analysis was treated as a grey box, in that some physics were imposed via the mean function. Simplifying the analysis to include a single mode in the band of interest was a reasonable assumption, as most of the physics (and variation), were accounted for with the SDOF mean function, and the remaining distortion was resolved via the black box component of the GP. (Note that the decision to focus on a narrow frequency band was made to simplify the analysis while demonstrating the proposed technology, although the form could be developed over a larger band and with a multiple DOF assumption).

To fit the real part, the real part of the FRFs from the fixed-free tests from each blade was copied 20 times, and random Gaussian noise with a magnitude equal to 5% of the absolute peak value of each FRF was added to each copy. The data copies were then concatenated into a vector, and 300 training points were randomly selected from the vector. The modal damping FRF model from Eq. (8) was used to generate a mean function, and the modal damping, residue, and natural frequency were used as the mean-function hyperparameters. Upper and lower bounds were applied to constrain the hyperparameter optimisation, based on prior knowledge of




Fig. 6. Predictions for supervised mixture of GPs using (a) real and (b) imaginary FRF data.

the blades and visual inspection of the data (e.g., the natural frequency was bounded between 40 and 60 Hz). In addition, the sign of the data (+/−) was assumed to be known *a priori*, and can easily be determined via data inspection. This process was repeated using the imaginary part of the FRFs, and was repeated for each blade. The posterior predictive mean of the GP is plotted with the mean function, variance, and training data for the real and imaginary FRFs in Figs. 6(a) and 6(b). For comparison, a single GP was learnt for the same data (although using a different training set comprised of 600 data points), as shown in Figs. 7(a) and 7(b).

Figs. 6(a) and 6(b) show that the mixture provided a reasonable fit for the FRF data, and the majority of the training data were enclosed by the variance bounds. Furthermore, the NMSE was calculated using Eq. (38) using an independent test data set ($M = 12480$), and showed acceptable values (2.4 for the real fit, and 6.9 for the imaginary fit). Conversely, Figs. 7(a) and 7(b) show that the differences in the natural frequencies of the blades resulted in significant variability in the data along the horizontal axis, for which the single GP is not well-suited. Because of this variation along the horizontal axis, the GP was forced to apply excessive damping when fitting the data. The NMSE was again calculated for the single GP ($M = 13080$), and was unsurprisingly, much higher than those for the mixture, at 55.8 for the real fit and 65.1 for the imaginary fit. These results suggest that it is more appropriate to use a mixture of GPs in cases where the data have significant horizontal variability, as occurred here.

4.2. A form using an unsupervised mixture of Gaussian processes (OMGP)

An overlapping mixture of Gaussian processes (OMGP), assumes that the number of trajectory functions are known *a priori*, but does not assume which data belong to a given class. In this case, because four helicopter blades contributed to the data set, it was assumed that four trajectories (i.e., $K = 4$ FRFs) describe the population form. The data were then classified and fitted by approximating the intractable posterior distribution $p(\{\mathbf{f}^{(k)}\}_{k=1}^K, \mathbf{Z} \mid \mathbf{x}, \mathbf{y})$ via the OMGP approach described in [13], using the variational inference and EM scheme laid out in Eqs. (18)–(37).

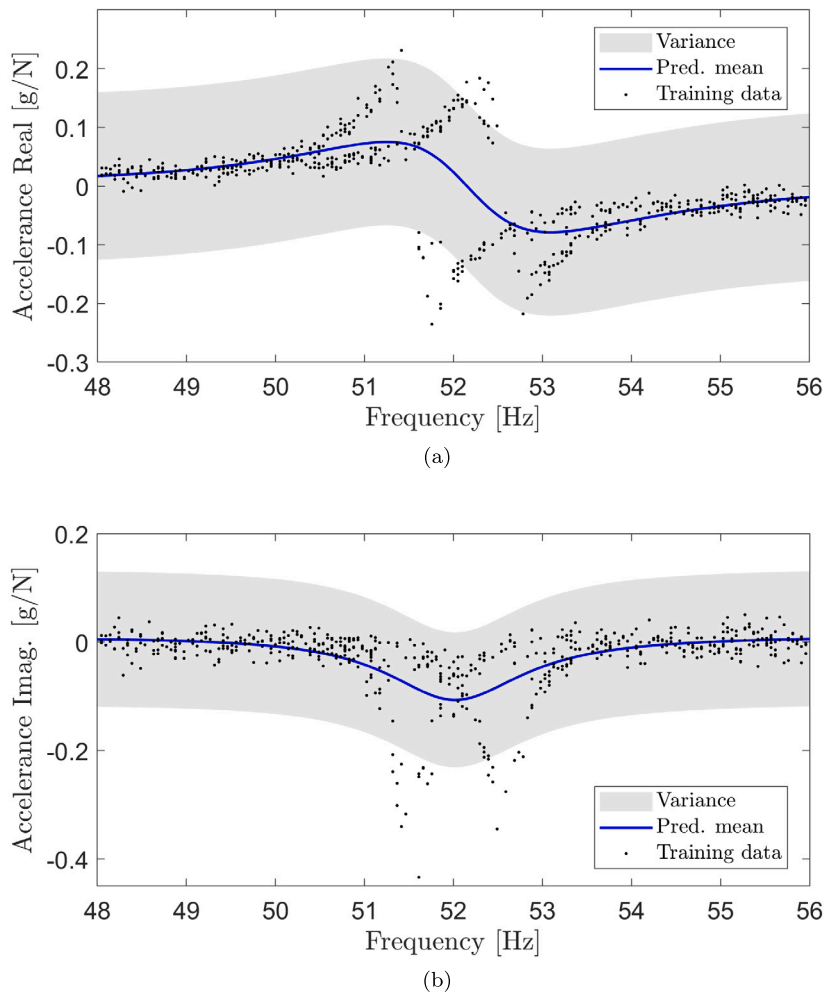


Fig. 7. Predictions for single GP using (a) real and (b) imaginary FRF data.

As with the previous example, a narrow frequency band containing the fifth bending mode of the blades, located between 48 and 56 Hz, was isolated for analysis and treated as an SDOF. Real and imaginary parts of the FRF were fit separately via Eqs. (18)–(37). The same replicated data with added noise were used as in the previous case, and 600 training points were randomly selected from these data. Likewise, the modal damping FRF model from Eq. (8) was used to generate a mean function, and modal damping, residue, and natural frequency were used as the mean-function hyperparameters. Upper and lower bounds were applied to constrain the hyperparameter optimisation, based on prior knowledge of the blades and visual inspection of the data, and the FRF direction was again assumed known. Because the real part of the FRF was more separable than the imaginary part (indeed, visual inspection of the FRFs showed that four distinct trajectories in the training data were harder to decipher from the imaginary part versus the real part), the OMGP performed much better for a variety of initialisations on the real data compared to the imaginary data. As such, the real data were fitted first, and an optimal solution was found by randomly initialising the process (i.e., via prior guesses for the hyperparameters), ten times while keeping the training set fixed, and choosing the solution with the maximum lower bound and therefore the highest marginal likelihood. The initialisation parameters were chosen randomly, but were restricted to reasonable intervals in the same manner as the hyperparameters were bounded during the optimisation. Once learnt, the optimised hyperparameters for the real data were used to initialise the OMGP for the imaginary data. The posterior predictive mean of the OMGP is plotted with the mean function, variance, and training data for the real and imaginary FRFs in Figs. 8(a) and 8(b).

Visual inspection of Figs. 8(a) and 8(b) show that the OMGP provided comparable results to the supervised mixture from the previous example (shown in Figs. 6(a) and 6(b)). Likewise, the NMSE were again calculated via Eq. (38) using an independent test data set ($M = 13080$). The calculated NMSE (2.3 for the real fit and 6.5 for the imaginary fit) were similar to those obtained with the supervised mixture (2.4 for the real fit, and 6.9 for the imaginary fit). The NMSE and MSD (computed via Eq. (39)) for the OMGP are listed in Table 2, along with those computed for the supervised mixture and single GP in the previous example.

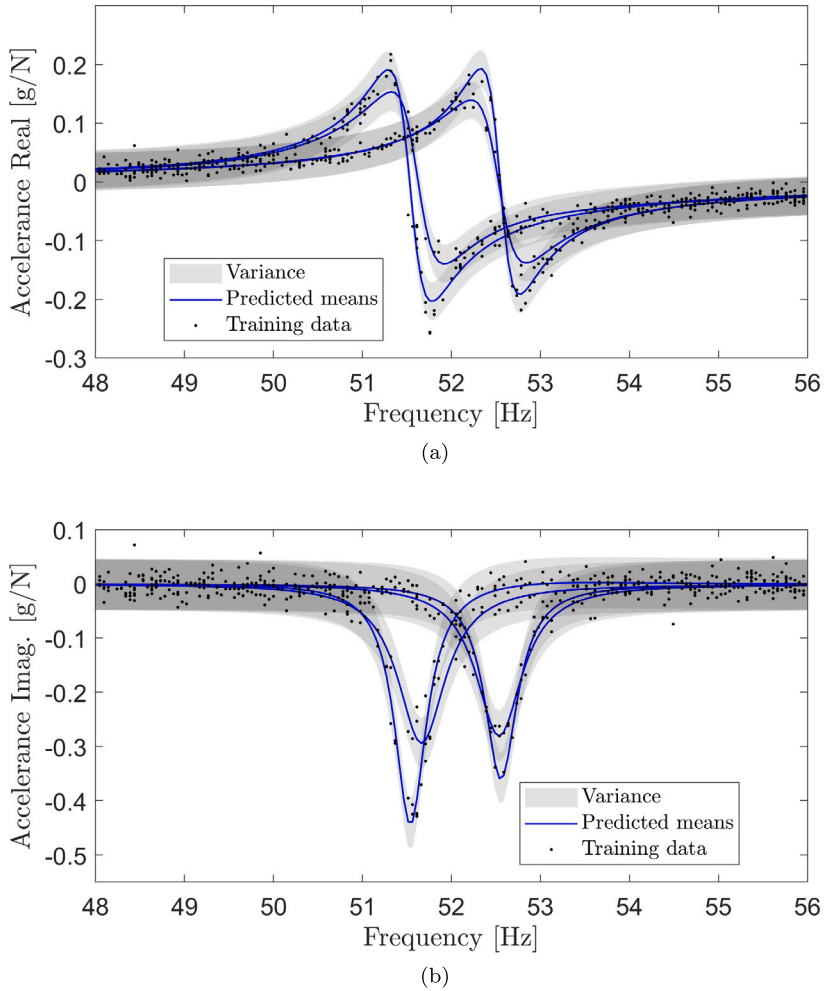


Fig. 8. OMGP predictions using (a) real and (b) imaginary FRF data.

Table 2
Model performance metrics for (OM)GP forms.

	Single GP		S. mixture		OMGP	
	real	imag	real	imag	real	imag
NMSE	55.76	65.05	2.40	6.92	2.32	6.49
MSD	1.01	0.71	0.99	0.98	0.76	0.57

The OMGP provides a unified form that generalises the normal condition for a population of structures, despite slight changes in boundary conditions and differences among the nominally-identical members. Unlike the conventional GP, that cannot accommodate large amounts of input-dependent noise, the OMGP (and the supervised mixture of GPs) can account for the horizontal variance that results from normal variations. The OMGP further generalises from the supervised mixture of GPs by allowing for an unlabelled training data set.

5. Novelty detection via the form

One practical implementation of the OMGP form is for damage identification, where unseen data can be compared against the posterior predictive distributions of the OMGP and evaluated for novelty. This section demonstrates how novelty detection could be accomplished, using a technique similar to that involving the MSD in [3,4,12], but using the marginal likelihood from Eq. (36) as a novelty index. Because the real and imaginary parts of the FRFs were fitted separately, the sum of the (negative log) marginal likelihood from each fit provided the novelty index. A Monte Carlo approach (as in [12]) was used to determine an appropriate normal-condition threshold.

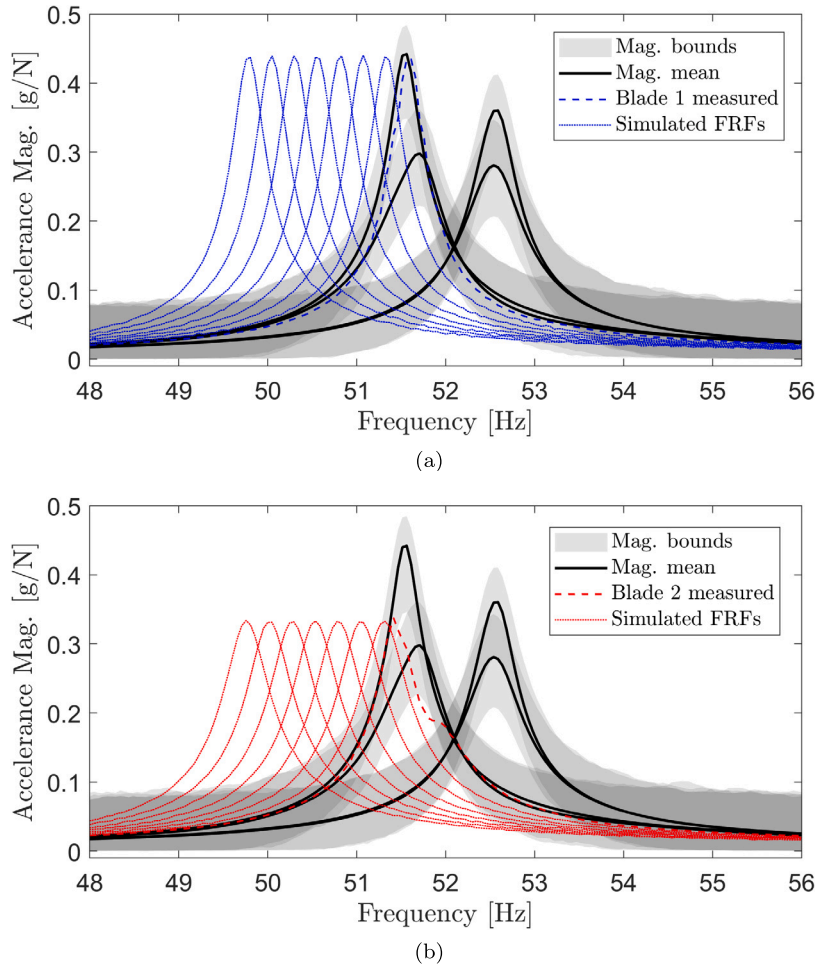


Fig. 9. Synthesised FRFs with incrementally decreasing natural frequency for (a) Blade 1 and (b) Blade 2, superimposed on the magnitude of the predicted OMGP.

To evaluate the sensitivity of the form to changes in natural frequency, real and imaginary SDOF FRFs were synthesised using Eq. (8) to simulate stiffness reduction resulting from damage, which may present as a downward shift in natural frequency. FRFs were generated for each blade using modal information from the fixed-free tests (i.e., residues and modal damping obtained via hyperparameter optimisation from fitting the OMGP prior mean functions — representing the ‘true’ values for these parameters) and natural frequency incrementally decreased by 0.5% for a maximum reduction of 3.5% (i.e., the natural frequency in the band of interest associated with each blade was reduced such that it was 0.5%, 1%, 1.5%, 2%, 2.5%, 3%, and 3.5% lower in frequency than the true, undamaged value). In other words, the ‘damaged’ FRFs were very similar to the undamaged experimental FRFs, except that they were slightly lower in frequency. The simulated FRFs were copied 1000 times and random Gaussian noise with a magnitude equal to 5% of the absolute peak value of the FRF was added to each copy. The negative log-marginal-likelihood was then computed for each copy, with the novelty index equal to the negative log sum of the marginal likelihoods for the real and imaginary parts. Likewise, the experimental data from the fixed-free tests for each blade were copied 1000 times and random Gaussian noise with a magnitude equal to 5% of the absolute peak value of the FRF was added to each copy. As with the simulated data, the copied experimental data (with added noise), were tested against the OMGP form, to evaluate the performance of the novelty detector for data known to be inlying.

To identify a suitable normal-condition threshold, a similar technique to that applied in [3,4,12] was used, where 1000 samples were randomly selected from the normal-condition data (i.e., experimental FRFs, replicated with added noise) used to train the form. The negative log-marginal-likelihood was then calculated for the samples, for a large number of trials. The critical value was the threshold with a certain percentage of the calculated values below it. This work used a 99% confidence interval, or 2.58σ .

An approximation of the FRF form magnitude was computed for visualisation purposes. The posterior predictive distributions of the real and imaginary OMGPs were sampled 10,000 times, and the FRF magnitudes were computed from these samples. The mean and uncertainty bounds of the resulting magnitudes are shown in Figs. 9(a)–10(b), along with the magnitudes of the simulated FRFs.




Fig. 10. Synthesised FRFs with incrementally decreasing natural frequency for (a) Blade 3 and (b) Blade 4, superimposed on the magnitude of the predicted OMGP.

Note that in Figs. 9(a)–10(b), ‘Mag. mean’ refers to the mean of the FRF magnitudes as computed from samples from the OMGP posterior while ‘Mag. bounds’ shows the range within the magnitudes. The negative log-marginal-likelihood computed using the experimental data from the fixed–free tests as test data and the various synthesised FRFs as test data for Blades 1–4 are shown in Figs. 11(a)–11(d), respectively.

Examining the negative log-marginal-likelihood from right to left to correspond with the downward frequency shift of the FRFs, Fig. 11(a) shows that $-\log p(y_* | x_*, D)$ computed using replicated experimental FRFs for Blade 1 approached but did not surpass the threshold (shown in Figs. 11(a)–11(d) as a solid red line), which is expected, given that the peak belonging to Blade 1 (as shown in Fig. 9(a)) was located at the edge of the training data with respect to frequency. As the natural frequency of the synthesised FRFs decreased, the novelty index became increasingly outlying. In Fig. 11(b), similar behaviour for Blade 2 is visible, although the novelty index was further below the threshold for each case as the FRF from Blade 2 had a lower magnitude than that of Blade 1, as shown in Figs. 9(a) and 9(b). Fig. 11(c) shows that for Blade 3, the novelty index gradually increased until approaching the threshold at -1.5% , which corresponded to the space between the peak groupings. As the natural frequency of the synthesised FRFs for Blade 3 decreased further (-1.5% to -2%), the novelty index fell, because the simulated FRFs aligned with the natural frequencies of Blades 1 and 2. As the natural frequency of the synthesised FRFs continued to decrease (-2.5% to -3.5%), the novelty index became increasingly outlying. Fig. 11(d) shows similar results for Blade 4; however, because Blade 4 had a larger FRF magnitude than Blade 3, the novelty index was closer to the threshold. These results suggest that for Blades 1 and to a lesser extent Blade 3, the technique would be quite sensitive to downward frequency shifts. The technique would be less sensitive to frequency reductions for Blades 3 and 4, as a decrease in frequency may result in overlap with the normal-condition peaks for Blades 1 and 2. Note that this metric was computed in a functional sense, meaning that the test data included the full (band-limited) FRF, rather than individual test points, and the full covariance matrices of the OMGP were used to compute the novelty index. Considering the metric in a functional sense means that if the test data fit a function significantly different than the form (e.g., the test data were not an FRF),

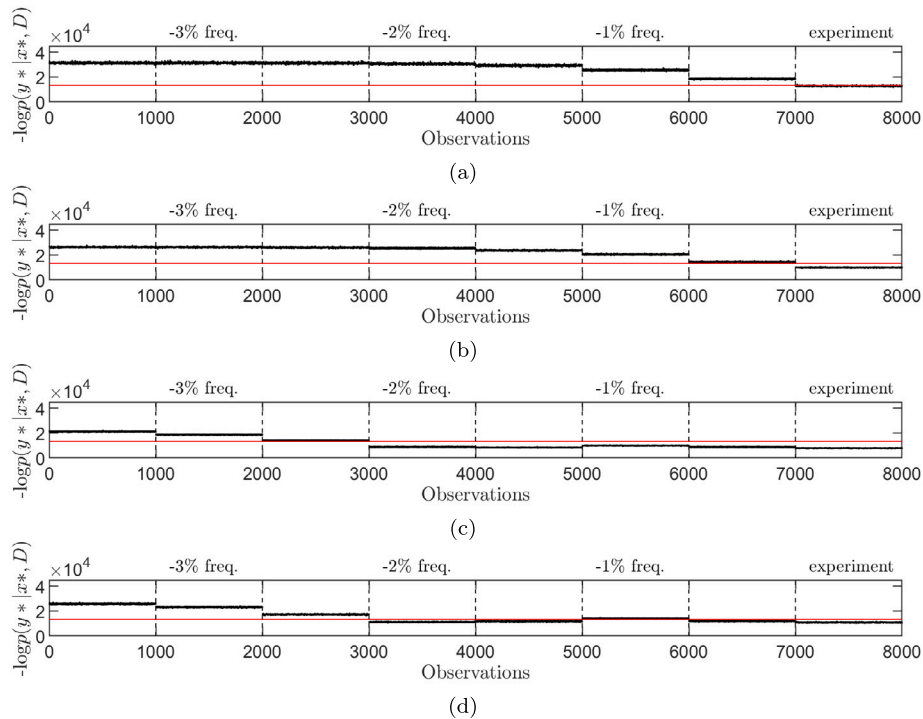


Fig. 11. Negative log-marginal-likelihood computed using experimental and simulated damaged results as test data for (a) Blade 1, (b) Blade 2, (c) Blade 3, and (d) Blade 4.

then the data could be flagged as outlying regardless of whether the majority of individual test points were within the variance bounds.

6. Concluding remarks

This paper presented data obtained from a comprehensive testing campaign on four healthy, nominally-identical helicopter blades, and used the measured data to develop a generic model, called a *form*, to represent the small ‘population’. In the first example presented, a supervised mixture of GPs was used to develop the form, where the labels were assumed known *a priori*. The second example used an (unsupervised) OMGP to develop the form, where the data classes were presumed unknown and needed to be inferred via the OMGP (variational inference/EM scheme) approach. Comparable results were obtained from both techniques. In addition, the posterior predictive distributions from the OMGP were applied to a novelty detection approach to evaluate new (simulated) data, providing insight into the generalisation capabilities of the model.

Population-based SHM seeks to transfer valuable information across similar structures, such as normal conditions and damaged states. The results presented herein demonstrate preliminary techniques for establishing a normal condition for a population of structures that are nominally-identical, but distinct, with consideration for variations such as discrepancies in material properties or fluctuations in boundary conditions. Although the proposed model provides an accurate representation of the population data, it does not (in its current form) consider correlation between the related functions. Further work will consider methods of inferring correlated task parameters for knowledge transfer within the group.

CRedit authorship contribution statement

T.A. Dardeno: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **L.A. Bull:** Conceptualization, Methodology, Software, Validation, Writing – review & editing. **R.S. Mills:** Investigation, Validation, Resources. **N. Dervilis:** Conceptualization, Methodology, Validation, Resources, Writing – review & editing. **K. Worden:** Conceptualization, Methodology, Validation, Resources, Writing – review & editing, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

The authors gratefully acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC), via grant references EP/R003645/1 and EP/R004900/1. L.A. Bull was supported by Wave 1 of the UKRI Strategic Priorities Fund under the EPSRC grant EP/W006022/1, particularly the *Ecosystems of Digital Twins* theme within that grant and the Alan Turing Institute. For the purpose of open access, the author(s) has/have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

This research made use of The Laboratory for Verification and Validation (LVV), which was funded by the EPSRC (grant numbers EP/J013714/1 and EP/N010884/1), the European Regional Development Fund (ERDF) and the University of Sheffield. The authors would like to extend special thanks to Michael Dutchman, for assisting with experimental setup.

References

- [1] A. Dowsett, D. O'Boy, S. Walsh, A. Abolfathi, S. Fisher, The prediction of measurement variability in an automotive application by the use of a coherence formulation, *Proc. Inst. Mech. Eng. D: J. Automob. Eng.* 232 (12) (2018) 1694–1700, <http://dx.doi.org/10.1177/0954407017734768>.
- [2] C. Chen, D. Duhamel, C. Soize, Probabilistic approach for model and data uncertainties and its experimental identification in structural dynamics: Case of composite sandwich panels, *J. Sound Vib.* 294 (1) (2006) 64–81, <http://dx.doi.org/10.1016/j.jsv.2005.10.013>.
- [3] L.A. Bull, P.A. Gardner, J. Gosliga, T.J. Rogers, N. Dervilis, E.J. Cross, E. Papatheou, A.E. Maquire, C. Campos, K. Worden, Foundations of population-based SHM, Part I: Homogeneous populations and forms, *Mech. Syst. Signal Process.* 148 (2021) 107141, <http://dx.doi.org/10.1016/j.ymssp.2020.107141>.
- [4] L.A. Bull, T.J. Rogers, N. Dervilis, E.J. Cross, K. Worden, A Gaussian process form for population-based structural health monitoring, in: *Proceedings of the 13th International Conference on Damage Assessment of Structures (DAMAS 2019)*, Porto, Portugal, Springer, 2019.
- [5] J. Gosliga, P.A. Gardner, L.A. Bull, N. Dervilis, K. Worden, Foundations of population-based SHM, Part II: Heterogeneous populations—graphs, networks, and communities, *Mech. Syst. Signal Process.* 148 (2021) 107144, <http://dx.doi.org/10.1016/j.ymssp.2020.107144>.
- [6] P.A. Gardner, L.A. Bull, J. Gosliga, N. Dervilis, K. Worden, Foundations of population-based SHM, Part III: Heterogeneous populations—mapping and transfer, *Mech. Syst. Signal Process.* 149 (2021) 107142, <http://dx.doi.org/10.1016/j.ymssp.2020.107142>.
- [7] K. Worden, H. Sohn, C. Farrar, Novelty detection in a changing environment: Regression and interpolation approaches, *J. Sound Vib.* 258 (2002) 741–761, <http://dx.doi.org/10.1006/jsvi.2002.5148>.
- [8] S. Alampalli, Influence of in-service environment on modal parameters, in: *Proceedings of the 16th International Modal Analysis Conference, IMAC, Society for Experimental Mechanics, Santa Barbara, CA, USA, 1998*.
- [9] P. Cawley, Long range inspection of structures using low frequency ultrasound, in: *Proceedings of the International Conference on Damage Assessment of Structures (DAMAS 97)*, Sheffield, UK, Sheffield Academic Press, 1997, pp. 1–17.
- [10] H. Sohn, Effects of environmental and operational variability on structural health monitoring, *Philos. Trans. A: Math. Phys. Eng. Sci.* 365 (2007) 539–560, <http://dx.doi.org/10.1098/rsta.2006.1935>.
- [11] L.A. Bull, P.A. Gardner, T.J. Rogers, N. Dervilis, E.J. Cross, E. Papatheou, A.E. Maquire, C. Campos, K. Worden, Bayesian modelling of multivalued power curves from an operational wind farm, *Mech. Syst. Signal Process.* (2021) 108530, <http://dx.doi.org/10.1016/j.ymssp.2021.108530>.
- [12] K. Worden, G. Manson, N. Fieller, Damage detection using outlier analysis, *J. Sound Vib.* 229 (3) (2000) 647–667, <http://dx.doi.org/10.1006/JSVI.1999.2514>.
- [13] M. Lázaro-Gredilla, S. Van Vaerenbergh, N.D. Lawrence, Overlapping mixtures of Gaussian processes for the data association problem, *Pattern Recognit.* 45 (4) (2012) 1386–1395, <http://dx.doi.org/10.1016/j.patcog.2011.10.004>.
- [14] S.-K. Hong, B.I. Epureanu, M.P. Castanier, D.J. Gorsich, Parametric reduced-order models for predicting the vibration response of complex structures with component damage and uncertainties, *J. Sound Vib.* 330 (6) (2011) 1091–1110, <http://dx.doi.org/10.1016/j.jsv.2010.09.022>.
- [15] J. Yang, N. Kessissoglou, Modal analysis of structures with uncertainties using polynomial chaos expansion, in: *Proceedings of Acoustics, Australian Acoustical Society, Victor Harbor, Australia, 2013*.
- [16] E.J. Cross, S.J. Gibson, M.R. Jones, D.J. Pitchforth, S. Zhang, T.J. Rogers, Physics-informed machine learning for structural health monitoring, in: A. Cury, D. Ribeiro, F. Ubertini, M.D. Todd (Eds.), *Structural Health Monitoring Based on Data Science Techniques*, Springer International Publishing, 2022, pp. 347–367, <http://dx.doi.org/10.1007/978-3-030-81716-9>.
- [17] K. Worden, G.R. Tomlinson, *Nonlinearity in Structural Dynamics: Detection, Identification and Modelling*, Institute of Physics Publishing, Bristol, UK, 2001.
- [18] K.P. Murphy, *Machine Learning: A Probabilistic Perspective*, The MIT Press, Cambridge, MA, USA, 2012.
- [19] C.E. Rasmussen, C.K.I. Williams, *Gaussian Processes for Machine Learning*, in: *Adaptive Computation and Machine Learning*, The MIT Press, Cambridge, MA, USA, 2006, p. 248.
- [20] M.C.K. Tay, C. Laugier, Modelling smooth paths using Gaussian processes, in: *Proceedings of the International Conference on Field and Service Robotics, FSR07*, Springer, Chamonix, France, 2007, pp. 381–390.
- [21] D.M. Blei, A. Kucukelbir, J.D. McAuliffe, Variational inference: A review for statisticians, *J. Amer. Statist. Assoc.* 112 (518) (2017) 859–877.
- [22] D.J.C. MacKay, *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, Cambridge, UK, 2003.
- [23] N.J. King, N.D. Lawrence, Fast variational inference for Gaussian process models through KL-correction, in: *Proceedings of the 17th European Conference on Machine Learning, EMCL 06*, Springer-Verlag, Berlin, Germany, 2006, pp. 270–281.
- [24] M. Lázaro-Gredilla, M.K. Titsias, Variational heteroscedastic Gaussian process regression, in: *Proceedings of the 28th International Conference on Machine Learning, ICML 11*, Omnipress, Bellevue, WA, USA, 2011.
- [25] J.S. Bendat, A.G. Piersol, *Random Data: Analysis and Measurement Procedures*, 4th Ed., John Wiley & Sons, New York, NY, 2011.
- [26] T.A. Dardeno, L.A. Bull, N. Dervilis, K. Worden, Investigating experimental repeatability and feature consistency in vibration-based SHM, in: *Proceedings of the 13th International Workshop on Structural Health Monitoring, IWSHM, 2022*.
- [27] T.A. Dardeno, M. Haywood-Alexander, R.S. Mills, L.A. Bull, N. Dervilis, K. Worden, Investigating the effects of ambient temperature on feature consistency in vibration-based SHM, in: *Proceedings of the 13th International Workshop on Structural Health Monitoring, IWSHM, 2022*.