



Domain-adapted Gaussian mixture models for population-based structural health monitoring

Paul Gardner¹ · Lawrence A. Bull² · Nikolaos Dervilis¹ · Keith Worden¹

Received: 3 November 2021 / Revised: 7 March 2022 / Accepted: 8 March 2022
© The Author(s) 2022

Abstract

Transfer learning, in the form of domain adaptation, seeks to overcome challenges associated with a lack of available health-state data for a structure, which severely limits the effectiveness of conventional machine learning approaches to structural health monitoring (SHM). These technologies utilise labelled information across a population of structures (and physics-based models), such that inferences are improved, either for the complete population, or for particular target structures — enabling a population-based view of SHM. The aim of these methods is to infer a mapping between each member of the population's feature space (called a domain) in which a classifier trained on one member of the population will generalise to the remaining structures. This paper introduces the domain-adapted Gaussian mixture model (DA-GMM) for population-based SHM (PBSHM) scenarios. The DA-GMM, infers a linear mapping that transforms target data from one structure onto a Gaussian mixture model that has been inferred from source data (from another structure). The proposed model is solved via an expectation maximisation technique. The method is demonstrated on three case studies: an artificial dataset demonstrating the approach's effectiveness when the target domain differs by two-dimensional rotations; a population of two numerical shear-building structures; and a heterogeneous population of two bridges, the Z24 and KW51 bridges. In each case study, the method is shown to provide informative results, outperforming other conventional forms of GMM (where no target labelled data are assumed available), and provide mappings that allow the effective exchange of labelled information from source to target datasets.

Keywords Population-based structural health monitoring · Domain adaptation · Domain-adapted Gaussian mixture model · Transfer learning

Lawrence A. Bull, Nikolaos Dervilis and Keith Worden have contributed equally to this work.

✉ Paul Gardner
p.gardner@sheffield.ac.uk

Lawrence A. Bull
lbull@turing.ac.uk

Nikolaos Dervilis
n.dervilis@sheffield.ac.uk

Keith Worden
k.worden@sheffield.ac.uk

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

² The Alan Turing Institute, The British Library, London NW1 2DB, UK

1 Introduction

Data-based approaches to structural health monitoring (SHM) are often limited by the range of available labelled health-state data in training. This lack of available health-state data means that most approaches are limited to novelty detection, or only being able to identify previously-seen health-states. Population-based structural health monitoring (PBSHM) is a branch of SHM that seeks to overcome these challenges by expanding the set of available labelled health-state data by considering a population of structures [1–4]. By pooling together datasets from a population, there is likely to be more labelled information available and an increase in the potential to diagnose more health-states of interest.

A significant challenge in adopting PBSHM is that the feature spaces between structures will not be aligned; this will occur for several reasons. First, the structures themselves can be different, whether because of manufacturing variations in homogeneous populations (i.e. nominally-identical structures), or because the

population is heterogeneous [1–4]. Second, issues arise because the monitoring setup is not exactly equivalent for each member of the population. Third, the environmental and operational conditions for each member of the population are different, meaning each member of the population may be observing a different part of the feature space. These shifts in the feature space (also known as domain shift), mean that labelled data cannot be naïvely shared by directly applying a classifier trained on one member of the population to other members. Instead, a mapping must be inferred between the feature spaces such that a classifier can be inferred on a harmonised dataset that will generalise between all members of the population.

Transfer learning, and more specifically domain adaptation, is a branch of machine learning that aims to achieve this goal of identifying a mapping between different domains (i.e. different feature spaces) [5–7]. By inferring a mapping that harmonises these domains based on some criteria (typically using statistical distances [8–14] and/or manifold assumptions [12, 15]), a classifier can be inferred for a source domain where labelled information is known, and be applied to unlabelled target domains. Several domain adaptation methods have been used within the SHM literature, with most focussing on identifying a latent space where the two feature datasets are harmonised [8–17]. The majority of the approaches are deterministic in nature, and utilise a distance metric to infer the mapping. This paper takes a different viewpoint, instead seeking to infer a mapping from the target domain directly onto the source domain (rather than to a latent space). In addition, the method outlined in this paper is probabilistic and uses a maximum likelihood criteria to infer the mapping.

The method outlined in this paper— named the domain-adapted Gaussian mixture model (DA-GMM)— is an extension of work proposed by Paaßen et al. [18], that inferred a linear mapping between a fully-labelled target dataset and a source Gaussian Mixture Model. The novelty in this paper is that the method has been extended to the scenario where the target is unlabelled (and even the scenario where the source is unlabelled as well). This extension makes the approach practical for PBSHM scenarios, with the method demonstrated on three case studies: an artificial dataset, a population of two numerical shear-building structures, and the Z24 [19] and KW51 [20] bridge datasets. A MATLAB implementation accompanies this paper—<https://github.com/pagard/EngineeringTransferLearning>.

The outline of this paper is as follows. Sect. 2 introduces the DA-GMM and the expectation maximisation algorithm used to solve the model. An artificial dataset is introduced in Sect. 3, to demonstrate the flexibility of the linear mapping in the context of two-dimensional rigid rotations. A case study inferring the mapping between a population of two numerical shear-building structures is presented in Sect. 4, where the approach is benchmarked against several other GMM models. In Sect. 5, the method is applied in a completely unsupervised manner between two bridges, the Z24 and KW51. Finally, conclusions are drawn in Sect 6.

2 Domain-adapted-Gaussian mixture model

The domain-adapted Gaussian mixture model (DA-GMM), an extension of the work by Paaßen et al. [18], seeks to learn a mapping from a target dataset onto a finite Gaussian Mixture Model (GMM) learnt in a supervised manner from a labelled source domain dataset $\{X_s, y_s\}$ (where $X_s \in \mathbb{R}^{D \times N_s}$ are the feature data and $y_s \in \mathbb{R}^{1 \times N_s}$ are the corresponding labels¹). Specifically, the approach seeks to learn a linear mapping via a projection matrix $H \in \mathbb{R}^{D \times D}$ that transforms the distribution of the unlabelled target feature data $X_t \in \mathbb{R}^{D \times N_t}$ such that the likelihood of the transformed data $\hat{X}_t = HX_t \in \mathbb{R}^{D \times N_t}$ being generated by the k -component source GMM is maximised. The underlying distribution of the transformed target data is, therefore, assumed to be defined by a k -component mixture model, where each of the k components is defined by a Gaussian distribution,

$$p(\hat{x}_t^{(i)} | z^{(i)} = k) = \mathcal{N}(Hx_t^{(i)} | \mu_s^{(k)}, \Sigma_s^{(k)}), \quad (1)$$

where k indexes the class group $k \in \{1, \dots, K\}$, where the latent variable $z^{(i)} \in \{1, \dots, K\}$ represents the mixture component for the feature data $\hat{x}_t^{(i)}$, and the mean $\mu_s^{(k)}$ and covariance $\Sigma_s^{(k)}$ of each component are defined by the supervised source domain GMM, with the source GMM being inferred from the labelled dataset $\{X_s, y_s\}$, (and hence fixed in the following inference).

The latent variable that governs the mixture $z^{(i)} \in \{1, \dots, K\}$ is categorically distributed such that,

$$P(z^{(i)} = k) = \pi^{(k)} \quad (2)$$

and $\sum_{k=1}^K \pi^{(k)} = 1$, where $\pi = \{\pi^{(1)}, \dots, \pi^{(k)}\}$ define the mixing proportions.

The posterior probability of each class, given a transformed data point is,

$$p(z^{(i)} = k | \hat{x}_t^{(i)}) = \frac{p(\hat{x}_t^{(i)} | z^{(i)} = k)P(z^{(i)} = k)}{\sum_{k=1}^K p(\hat{x}_t^{(i)} | z^{(i)} = k)P(z^{(i)} = k)}, \quad (3)$$

$$p(z^{(i)} = k | \hat{x}_t^{(i)}) = \frac{\pi^{(k)} \mathcal{N}(Hx_t^{(i)} | \mu_s^{(k)}, \Sigma_s^{(k)})}{\sum_{k=1}^K \pi^{(k)} \mathcal{N}(Hx_t^{(i)} | \mu_s^{(k)}, \Sigma_s^{(k)})}, \quad (4)$$

where $p(z^{(i)} = k | Hx_t^{(i)}) = r^{(i,k)}$, aptly named the *responsibility matrix*. A maximum likelihood estimate of the labels \hat{y}_t can be obtained by selecting the class with the highest probability in the responsibility matrix for each observation,

¹ These labels provide insight into the physical phenomena represented by the feature data, for example, a label may indicate that data points correspond to a damage state, which may indicate location on the structure.

$$\hat{y}_t^{(i)} = \arg \max_{k \in \{1, \dots, K\}} [p(z^{(i)} = k | H\mathbf{x}_t^{(i)})]. \tag{5}$$

Parameter inference can be performed for the projection matrix H , as well as the mixing proportions for the transformed target data $\boldsymbol{\pi}$, using a maximum likelihood approach. The complete data (i.e. \hat{X}_t and Z combined) log likelihood of the projected target data given the parameter set $\boldsymbol{\theta} = \{H, \boldsymbol{\pi}\}$ is defined as,

$$\log p(\hat{X}_t, Z | \boldsymbol{\theta}) = \sum_{n=1}^N \log \sum_{k=1}^K \pi^{(k)} \mathcal{N}(H\mathbf{x}_t^{(i)} | \boldsymbol{\mu}_s^{(k)}, \Sigma_s^{(k)}), \tag{6}$$

which can be maximised using Expectation Maximisation (EM).

EM is a maximum likelihood technique for performing inference when missing values causes the marginal likelihood to be intractable [21], as is the case with the DA-GMM, where the latent variables Z are unobserved and their distribution is unknown. Instead, EM iterates between an expectation step (E-step) and a maximisation step (M-step), to find a maximum likelihood estimate of the marginal likelihood. The E-step requires the definition of the auxiliary function $Q(\boldsymbol{\theta}, \boldsymbol{\theta}_0)$, which, using the current parameter estimates $\boldsymbol{\theta}_0$, is the expectation of the complete log likelihood function with respect to the current posterior distribution of the latent variables. The M-step then finds a new set of parameter estimates $\boldsymbol{\theta}$ by maximising the auxiliary function with respect to the parameters. The steps are repeated until convergence criteria are met [21].

The auxiliary function (the expected complete data log likelihood) for the DA-GMM is,

$$Q(\boldsymbol{\theta}, \boldsymbol{\theta}_0) = \mathbb{E}_{Z|\hat{X}_t, \boldsymbol{\theta}_0} [\log p(\hat{X}_t, Z | \boldsymbol{\theta})], \tag{7}$$

where the expectation can be rewritten as,

$$Q(\boldsymbol{\theta}, \boldsymbol{\theta}_0) = \sum_Z p(Z | \hat{X}_t, \boldsymbol{\theta}_0) \log p(\hat{X}_t, Z | \boldsymbol{\theta}), \tag{8}$$

which in turn can be simplified using an indicator function $\mathcal{I}(\cdot)$ for the mixing component for the i^{th} data point and substituting Eq. (6),

$$Q(\boldsymbol{\theta}, \boldsymbol{\theta}_0) = \sum_{i=1}^N p(z^{(i)} | \hat{\mathbf{x}}_t^{(i)}, \boldsymbol{\theta}_0) \sum_{k=1}^K \mathcal{I}(z^{(i)} = k) \log [\pi^{(k)} \mathcal{N}(H\mathbf{x}_t^{(i)} | \boldsymbol{\mu}_s^{(k)}, \Sigma_s^{(k)})], \tag{9}$$

where $p(z^{(i)} | \hat{\mathbf{x}}_t^{(i)}, \boldsymbol{\theta}_0) \mathcal{I}(z^{(i)} = k) = p(z^{(i)} = k | \hat{\mathbf{x}}_t^{(i)}) = r^{(i,k)}$. This expression simplifies to,

$$Q(\boldsymbol{\theta}, \boldsymbol{\theta}_0) = \sum_{i=1}^N \sum_{k=1}^K r^{(i,k)} (\log [\pi^{(k)}] + \log [\mathcal{N}(H\mathbf{x}_t^{(i)} | \boldsymbol{\mu}_s^{(k)}, \Sigma_s^{(k)})]). \tag{10}$$

The auxiliary function can be maximised with respect to the mixing proportions $\boldsymbol{\pi}$ analytically, by adding a Lagrange multiplier (to ensure that $\sum_{k=1}^K \pi^{(k)} = 1$) and setting the partial derivative to zero,

$$\pi^{(k)} = \frac{1}{N} \sum_{i=1}^N r^{(i,k)}. \tag{11}$$

The maximisation step with respect to the projection matrix H , only requires maximising the part of the auxiliary function dependent on H ,

$$\max_H \sum_{i=1}^N \sum_{k=1}^K r^{(i,k)} \log [\mathcal{N}(H\mathbf{x}_t^{(i)} | \boldsymbol{\mu}_s^{(k)}, \Sigma_s^{(k)})] \tag{12}$$

which can be expanded to,

$$\begin{aligned} \max_H \sum_{i=1}^N \sum_{k=1}^K r^{(i,k)} \log [& (2\pi^{-\frac{n}{2}} |\Sigma_s^{(k)}|^{-\frac{1}{2}})] \\ & \times -\frac{1}{2} (H\mathbf{x}_t^{(i)} - \boldsymbol{\mu}_s^{(k)})^T (\Sigma_s^{(k)})^{-1} (H\mathbf{x}_t^{(i)} - \boldsymbol{\mu}_s^{(k)}), \end{aligned} \tag{13}$$

where the constants with respect to H can be omitted, leading to the minimisation of a weighted quadratic error,

$$\min_H \sum_{i=1}^N \sum_{k=1}^K r^{(i,k)} (H\mathbf{x}_t^{(i)} - \boldsymbol{\mu}_s^{(k)})^T (\Sigma_s^{(k)})^{-1} (H\mathbf{x}_t^{(i)} - \boldsymbol{\mu}_s^{(k)}) =: \mathbb{E}_Q(H). \tag{14}$$

It can be shown that $\mathbb{E}_Q(H)$ is convex [18], and the gradient is,

$$\nabla_H \mathbb{E}_Q(H) = 2 \sum_{k=1}^K (\Sigma_s^{(k)})^{-1} \sum_{i=1}^N r^{(i,k)} (H\mathbf{x}_t^{(i)} - \boldsymbol{\mu}_s^{(k)}) \mathbf{x}_t^{(i)T} \tag{15}$$

meaning that the maximisation step with respect to H can be solved using a gradient-based optimiser. In the case where $\Sigma_k = \Sigma, \forall k$, the projection matrix H can be found in closed form as,

$$H = MRX_t^T (X_t^T X_t)^{-1}, \tag{16}$$

where $M = [\boldsymbol{\mu}^{(1)}, \dots, \boldsymbol{\mu}^{(k)}] \in \mathbb{R}^{D \times K}$ and $R(i, k) = r^{i,k}$, where $R \in \mathbb{R}^{K \times N}$. This result is similar to a least-squares solution for a multiple-output linear regression, and, therefore, leads to the possibility of extending the mapping to being nonlinear and adding a prior over H .

The algorithm is shown in Algorithm 1. The linear projection matrix needs to be initialised either randomly, or given some prior knowledge about the mapping. The number of classes K is defined completely by the labels if the source GMM is supervised, or should be selected via a model selection approach if an unsupervised GMM is utilised.

Algorithm 1 DA-GMM

```

1:  $\{\mu_s^{(k)}, \Sigma_s^{(k)}\}_{k=1}^K = GMM(X_s, y_s, K)$ 
2:  $E \leftarrow \infty$ ,  $H \leftarrow H_0$ ,  $\pi^{(k)} \leftarrow 1/k \forall k \in K$ 

3: while true do

4:   E-step
5:   for  $i = 1 : N$  do
6:     for  $k = 1 : K$  do
7:        $r^{(i,k)} = \text{Eqn. 4}$ 
8:     end for
9:   end for

10:  M-step
11:  for  $k = 1 : K$  do
12:     $\pi^{(k)} = \text{Eqn. 11}$ 
13:  end for
14:   $H = \min_H \mathbb{E}_Q(H)$  ▷ Using Eqn. 14
15:   $E' = \mathbb{E}_Q(H)$ 

16:  Convergence check
17:  if  $|E - E'| < \epsilon$  then return  $H$ 
18:  end if
19:   $E \leftarrow E'$ 
20: end while

```

3 Artificial dataset

To demonstrate the capability of the linear transform H in the DA-GMM, an artificial dataset is utilised to evaluate the method's ability to infer linear rotation matrices. The dataset consists of four, two-dimensional Gaussian components, as stated in Table 1. The source dataset was obtained by directly sampling from these distributions $X_s = X$, whereas the target data were created by applying a rotation to the data obtained from these distributions, i.e., $X_t = HX$, where H is parametrised given some angle α , as,

$$H(\alpha) = \begin{bmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{bmatrix}. \quad (17)$$

The target test data are obtained from 1000 observations of each class transformed by the same rotation matrix (i.e. $N_t^{\text{test}} = 4000$).

Figure 1 states the testing accuracy on the target domain dataset against the rotation in degrees (where $\alpha = \{5, 10, \dots, 355\}$). The results shown correspond to the greatest accuracy from 25 random initialisations of the EM algorithm (with the training and testing data remaining the same for all 25 repeats), ensuring an optimal solution has been

Table 1 Gaussian components of the artificial dataset and number of training data points in the source and target domains (N_s and N_t , respectively). Gaussian distributions are parametrised by a mean μ and covariance Σ , $\mathcal{N}(\mu, \Sigma)$

k	μ	Σ	$N_s^{(k)}$	$N_t^{(k)}$	Label, y
1	$[-4.0 \ -1.0]^T$	$\begin{bmatrix} 0.60 & 0.10 \\ 0.10 & 0.60 \end{bmatrix}$	500	100	0
2	$[1.0 \ -0.5]^T$	$\begin{bmatrix} 0.40 & -0.30 \\ -0.30 & 0.40 \end{bmatrix}$	150	50	1
3	$[4.5 \ 0.2]^T$	$\begin{bmatrix} 0.48 & 0.28 \\ 0.28 & 0.48 \end{bmatrix}$	100	50	2
4	$[-6.0 \ 2.0]^T$	$\begin{bmatrix} 0.30 & 0.20 \\ 0.20 & 0.30 \end{bmatrix}$	75	50	3




Fig. 1 Comparison of testing accuracies (—) on the target dataset for different rotations of the artificial dataset

obtained. Despite the stochastic nature of the random initial conditions, the test accuracies are consistent with rotation (with all $\approx 100\%$). This is clear evidence that the DA-GMM approach is robust to rotations in two-dimensions. As an illustration of the approach, Fig. 2 depicts the source and target datasets, along with the transformed target dataset from the DA-GMM, for a rotation of 80° (i.e. $X_t = H(80)X$).

4 Case study: shear-building structures

Within the field of PBSHM a challenging problem is to identify mappings between heterogeneous populations. This difficulty arises because, as structures become increasingly different, their features spaces are more dissimilar, typically making the mapping more complex. This case study considers a population of two numerical four degree-of-freedom shear-building structures, constructed from different material properties (one aluminium and the other steel) and geometries.

The numerical simulations are obtained from four degree-of-freedom lumped-mass models, where the outputs of the model were four damped natural frequencies




Fig. 2 An example of the artificial dataset when $\alpha = 80^\circ$. Panels (a) and (b) present the source data X_s (\cdot), and the untransformed target data X_t (\triangleleft), respectively. Panels (c) and (d) show the DA-GMM predictions on the transformed target \hat{X}_t (\triangle) training and testing data, respectively (where Panel (c), includes the source data for reference). Panels (a), (c) and (d) include visualisations of the inferred source Gaussian mixture model clusters denoted by μ (+) and 2σ (---)

$\{\omega_1, \omega_2, \omega_3, \omega_4\}$ (calculated in a similar way to [3, 12]). The mass of each floor was parametrised by a length l_f , width w_f , thickness t_f and density ρ . The stiffness between each floor is formed via four rectangular cross-sectioned cantilever beams parametrised by length l_b , width w_b , thickness t_b and Young’s Modulus E . Damage is introduced to the beams via an open crack, using the stiffness-reduction model proposed by Christides and Barr [22]. For each structure, five damage scenarios were considered, based on a localisation problem; the undamaged condition $y = 0$, and an open crack on one of the beams at each of the four floors ($y = \{1, 2, 3, 4\}$ numbered from the ground floor upwards). For each location, the crack had a length of 5% of the beam width and was located at 10% of the way up the beam’s length.

To introduce variability, the material properties $\{\rho, E\}$ and damping coefficients c were defined as random variables, where each output observation was obtained by a random draw from an underlying distribution. The properties for the two structures are shown in Table 2. The labelled training data for Structure One consisted of 100 observations of $y = 0$, and 75 observations for each of the four damage classes $y = \{1, 2, 3, 4\}$ (i.e. $N_s = 400$). The unlabelled training data for Structure Two comprised 50 observations of $y = 0$ and 30 observations of each damage class $y = \{1, 2, 3, 4\}$ (i.e. $N_t = 170$). The level of class imbalance was chosen to reflect SHM problems, where typically more normal condition data ($y = 0$), are available than for each damage class individually. For both structures the test datasets were comprised of 1000 observations of each class ($N_s^{test} = N_t^{test} = 5000$).

The PBSHM scenario was to infer a mapping from Structure One to Structure Two, such that the labelled data in Structure One can aid in classifying the unlabelled data from Structure Two, i.e. Structure One is the source domain and Structure Two the target. To aid visualisation, the features in this analysis were the first two principal components of the four damped natural frequencies (calculated from the training data from each domain individually), i.e. $X_s \in \mathbb{R}^{2 \times N_s}$ and $X_t \in \mathbb{R}^{2 \times N_t}$, for structures One and Two, respectively. The feature spaces are visualised in Fig. 3, where it can be seen that the source and target domains are very different and there is a large amount of overlap between classes $y = \{1, 2, 4\}$ in the target domain. It is hoped that the

Table 2 Geometric and material properties of the shear-building structures (1 and 2). Gaussian distributions are parametrised by a mean μ and variance σ^2 , $\mathcal{N}(\mu, \sigma^2)$, and Gamma distributions by a shape κ and scale θ , $\mathcal{G}(\kappa, \theta)$

	$\{l_f, w_f, t_f\}$, mm	$\{l_b, w_b, t_b\}$, mm	ρ , kg/m ³	E , GPa	c , Ns/m
1	{350, 254, 25}	{200, 25, 6.25}	$\mathcal{N}(2800, 1)$	$\mathcal{N}(71, 1.2 \times 10^{-9})$	$\mathcal{G}(50, 0.1)$
2	{300, 250, 24}	{210, 24, 6.24}	$\mathcal{N}(8000, 5)$	$\mathcal{N}(210, 1.1 \times 10^{-9})$	$\mathcal{G}(50, 0.1)$




Fig. 3 Comparison of Gaussian mixture models and their predictions for the shear-building case study on the training data; $y = 0$ is the normal condition, and $y = 1, 2, 3, 4$ denote damage at each floor. Source data X_s (\cdot), target data X_t (\triangleleft) and transformed target data \hat{X}_t (\triangle) are depicted against the inferred Gaussian mixture model clusters denoted by μ (+) and 2Σ (—). Panel (a) shows the source model applied to the target data and Panel (b) displays the DA-GMM predictions. Panels (c) and (d) show the supervised and unsupervised Gaussian mixture models inferred from the target data (where the unsupervised labels are determined by the proximity of the inferred clusters to the supervised model)

Table 3 Training and testing accuracies for the different Gaussian mixture model scenarios. $GMM_a(X_b)$ denotes a supervised Gaussian mixture model trained on domain a and applied to dataset X_b , where the subscript un denotes an unsupervised mixture model. It is noted that the unsupervised model labels are determined by the proximity of the inferred clusters to the supervised model and the accuracies are used here as a reference measure

	$GMM_s(X_s)$	$GMM_s(X_t)$	$GMM_s(\hat{X}_t)$	$GMM_t^{un}(X_t)$	$GMM_t(X_t)$
Training	100.0%	55.9%	84.7%	70.0%	94.7%
Testing	100.0%	44.0%	81.5%	76.8%	95.6%

DA-GMM will aid both in labelling the target domain and improving separability between these classes.

Five scenarios were examined, demonstrating the comparative performance of the DA-GMM approach against conventional maximum likelihood GMMs, with their accuracies compared in Table 3. The first scenario (a prerequisite for the DA-GMM), was training the source domain supervised GMM, where a testing accuracy of 100% was achieved. The second analysis applied the source domain GMM directly to the untransformed target data, as seen in Figure 3a, where a testing classification accuracy of 44.0% was achieved. This result shows that there is a need for performing domain adaptation, as there is domain shift between the source and target datasets. The second scenario in Figure 3b is where the target data have been transformed using

the DA-GMM. The target data have been expanded to match the source GMM and as such have improved the testing accuracy to 81.5%. It is noted that the misclassification that has occurred is because of the overlap in the untransformed target data and its effect on the mapping.

The final two scenarios consider the target data alone (i.e. there is no attempt to transfer knowledge), with Figure 3c showing the results of the fully supervised GMM trained on a labelled target dataset, and Figure 3d displaying the unsupervised GMM² results (where classes have been assigned by the proximity of the inferred unsupervised clusters to the supervised model). These results show that with perfect information, a testing accuracy of 95.6% is achieved, and with no label information, the best accuracy one could achieve (assuming classes can be assigned with inspection knowledge) would be 76.8%. This evidences the fact that the DA-GMM provides better classification accuracy than an unsupervised model on the target domain (if labels could be assigned), and is a robust way of labelling the target domain from source domain observations.

5 Case study: Z24 and KW51 bridges

Inferring mappings between structures such that label information can be shared is an important aspect of PBSHM. By linking structures via mappings, any labelled data obtained for one structure can be applied to the rest of the population. In this section, domain adaptation is performed, such that a mapping is obtained in a heterogeneous population of two partially-labelled³ bridge datasets, from the Z24 [19] and KW51 bridges [20]. By learning a mapping between the two bridge datasets, any future label information obtained for one structure can be directly applied to the other. This is particularly important for managers of bridge infrastructure, as any observation of damage, from any member of the population, can be used in diagnosing that health-state for any bridge in the population (given that a mapping can be inferred).

5.1 Z24 and KW51 datasets

The Z24 bridge dataset has been well-studied within the literature, with numerous analyses managing to identify the key events in the dataset [23–32]. The Z24 bridge, located in the canton Bern near Solothurn, Switzerland, was a concrete

² The unsupervised GMM was trained using expectation-maximisation.

³ These datasets are said to be *partially-labelled*, as labels are only recorded for experimental damage introduced to the Z24 dataset [19] and for the retrofit condition in the KW51 dataset [20] and not for all data points, i.e. the various normal condition states.

highway bridge that was used as an experimental structure for an SHM campaign before its demolition in 1998, as part of the SIMCES project [19]. The monitoring campaign occurred during the year before demolition, in which a range of environmental data, as well as the acceleration response at 16 locations, were measured. The acceleration responses were processed using operational modal analysis (OMA), such that the first four natural frequencies of the structure were obtained. Damage was introduced incrementally into the bridge, with relatively small scale damage beginning on the 10th August 1998 — the pier was incrementally lowered by a few centimetres — with more substantial damage occurring after the 31st August 1998, beginning with the failure of the concrete hinge. For a complete description of the benchmark dataset, the reader is referred to [19]. It is noted that model-updating approaches [23, 26, 27] have been able to detect when the pier was lowered by between 80–95mm, which occurred on the 17–18th August. In contrast, several data-based approaches [24, 28–32] have detected the onset of the smallest damage introduced to the Z24 bridge, a lowering of the pier by 20 mm on the 10th August, with several even detecting when the installation equipment was brought onto the bridge (9th August).

The KW51 bridge is a steel bowstring railway bridge in Leuven, Belgium. A 15-month monitoring campaign occurred between 2018 and 2019, in which the acceleration response, the strains in the deck and gauges of the rails, the displacement at the bearings, as well as environmental data were all recorded [20]. The acceleration responses were processed using OMA to obtain the first 14 natural frequencies of the structure. During the monitoring campaign, every diagonal member was retrofit with a steel box to strengthen the design of the bridge, with details of the retrofit specified in [20]. The bridge condition before the retrofit was measured from 2nd October 2018 until the 15th May 2019, at which point the retrofitting process was carried out. The retrofit was completed on 27th September 2019. Novelty detection has been successfully performed on this dataset using robust principal component analysis (PCA) and linear regression methods [33]. For a complete overview of the dataset the reader is referred to [20].

Although the two bridges are very different in design (forming a heterogeneous population), there are similarities in their modal responses⁴. Despite different absolute values in natural frequencies, the first and third natural frequencies of the Z24 have correspondence with the tenth and twelfth natural frequencies of the KW51; both represent vertical bending modes

⁴ For knowledge transfer to be attempted, structures in a population should have some physical correspondence [1–3]; this is to reduce the likelihood of negative transfer [3], in which transfer learning degrades the performance of the classifier.




Fig. 4 Z24 dataset. Ambient temperature, first (ω_1) and third (ω_3) natural frequencies. Horizontal line indicates 0 °C and the vertical black line the training and testing data split, with the red solid line denoting the onset of damage




Fig. 5 KW51 dataset. Ambient temperature, tenth (ω_{10}) and twelfth (ω_{12}) natural frequencies. Horizontal line indicates 0 °C and the vertical black line the training and testing data split, with the red solid lines denoting the start and end of the retrofit period

of the deck, with effectively the same nodal and anti-nodal pattern (see [19] and [20] for a visualisation of the Z24 and KW51 mode shapes, respectively). In addition, both bridges undergo a stiffening effect caused by cold temperature conditions. Figures 4 and 5 show the two datasets against ambient temperature, where it can be seen that the below-freezing conditions lead to a ‘stiffening-effect’ on the natural frequencies; in the Z24, the asphalt stiffens, compared to freezing of the rail ballast in the KW51. It is noted that, although these low-temperature conditions correlate to a below-freezing condition, there are no ‘ground truth’ labels available for these conditions in either dataset. Given the similarities in modal response and the existence of at least two normal conditions (normal ambient and low temperature), the datasets form a transfer learning problem which can be solved using the DA-GMM.

The datasets were divided into training and testing data for the DA-GMM mapping. The first 99 days of monitoring data from the Z24 are used as training data, as this time period




Fig. 6 Statistic-aligned Z24 (source) X_s , and KW51 (target) X_t , features. Dashed line indicates data points used to statistically align the data, the solid black lines indicate training and testing data splits and the solid red lines denote the onset of damage and the retrofit states, respectively

covers both normal condition behaviours; the remaining data, including when damage occurred, are used to test the mapping. The KW51 training data were formed from the first 1500 data points covering the two temperature conditions, with the remaining pre-retrofit and post-retrofit data being used as testing data for the mapping. The training and testing divisions for both datasets are indicated on Figs. 4 and 5.

5.2 Statistic alignment and the unsupervised domain-adapted Gaussian mixture model

The aim of the analysis in this section was to find a mapping using the DA-GMM approach that aligns the data from these two bridges for the two environmental conditions in a completely unsupervised manner (as no ‘ground truth’ labels are known for either dataset for these conditions). Although completely unlabelled, the Z24 is considered as the source dataset (i.e. it forms the dataset where the GMM is inferred) and the KW51 is considered as the target domain (i.e. the KW51 is mapped onto the Z24 dataset). This change means that line one of Algorithm 1 is modified such that the source mean and covariances are inferred using an unsupervised GMM (i.e. no source labels y_s are required) using an EM maximum-likelihood approach.

Before applying the DA-GMM, pre-processing is performed to aid transfer, in a similar manner to how normalisation is performed as part of best practice in machine learning algorithms. The pre-processing performs statistic alignment [7], a process of aligning the lower-order statistics of the feature space. In this case, the statistics of the first 100 data points ($N = 100$) are aligned to the target domain, with the aim of centring both datasets around the normal condition i.e.,

$$X_t = \frac{\omega_t - \mu_{1:N}(\omega_t)}{\sigma_{1:N}(\omega_t)} \tag{18}$$

$$X_s = \left(\frac{\omega_s - \mu_{1:N}(\omega_s)}{\sigma_{1:N}(\omega_s)} \right) \sigma_{1:N}(X_t) + \mu_{1:N}(X_t), \tag{19}$$

where ω_s and ω_t are the set of natural frequencies for each bridge, $\mu_{1:N}(\cdot)$ and $\sigma_{1:N}(\cdot)$ find the mean and standard deviation of the first N data points of their arguments, and X_s and X_t are the source and target pre-processed features. The statistically-aligned features are presented in Fig. 6, where it can be seen that the offset between the source and target has been removed (removing the need for X_t to be augmented by a matrix of ones), and the problem has been simplified to one of learning a rotation matrix between the source and target datasets. The pre-processing steps mean that the DA-GMM algorithm can be modified, with the parameter set becoming $\theta = \{ \alpha, \pi \}$, where the linear projection matrix is parametrised as in Eq. (17) and α is the angle of rotation; this reduces the number of parameters that need to be inferred by the DA-GMM model.

Figures 4 and 5 indicate that there is class imbalance between the ambient and low temperature classes. This prior assumption, that there will be fewer data points in the low temperature class when compared to the ambient temperature class, can be encoded into the DA-GMM to learn an optimal mapping that equally prioritises both classes. In this case study, the maximisation step is modified such that Eq. (14) becomes,

$$\mathbb{E}_Q(H) = \min_H \sum_{k=1}^K w_k \sum_{i=1}^N r^{(i,k)} (Hx_t^{(i)} - \mu_s^{(k)})^T (\Sigma_s^{(k)})^{-1} (Hx_t^{(i)} - \mu_s^{(k)}), \tag{20}$$

where w_k are a prior imbalance weight for each class. In this case study, the values of the weight vector w are assigned based on the portion of the target training dataset below 0°C, where the ambient condition is down-weighted and the low temperature is up-weighted; $w = \{0.055, 0.945\}$ for the ambient and low temperature conditions, respectively.

5.3 Domain adaptation results

A two-class unsupervised DA-GMM was inferred on the Z24 (unlabelled source), and KW51 (unlabelled target), statistically-aligned feature data. The inferred projection and source GMM are presented in Fig. 7a, with the natural frequencies of the bridges labelled according the predictive classes of the GMM displayed in Fig. 7b and c. It can be seen in Fig. 7a, that the target data have been well-aligned with the source data, and that the inferred feature space retains physical meaning. The ambient normal condition is centred around the origin and can be viewed as the ‘baseline’ relative natural frequencies, with the low temperature effect causing an increase in the feature values, corresponding to stiffening behaviour. The GMM has

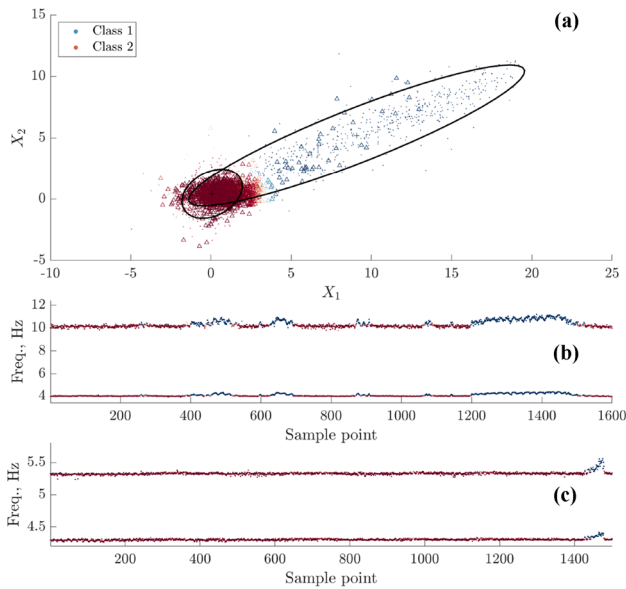


Fig. 7 Unsupervised domain-adapted Gaussian mixture model predictions of the Z24 and KW51 bridge datasets. Panel (a) is a comparison of the two features for each bridge against the probability of being in either class, with the inferred (source) unsupervised Gaussian mixture model for reference (μ (+) and 2Σ (—)). The Z24 is denoted X_s (·) and the KW51 \hat{X}_t (Δ). Panels (b) and (c) are the Z24 (ω_1 and ω_3) and KW51 (ω_{10} and ω_{12}) natural frequencies against sample point

identified the transition from the ambient to low temperature classes for both bridges, as shown in Fig. 7b and c, meaning that one GMM can be used to diagnose behaviour on both bridges. This result demonstrates the power of the DA-GMM, that two bridges of different design can be mapped onto a single feature space where classification can be performed for the complete population.

Once the mapping has been inferred, any future data can be projected onto the harmonised feature space. The testing data for the Z24 and KW51 were mapped into this space and an unsupervised GMM inferred using maximum-likelihood EM. Given that the number of classes is unknown, model selection was performed. The number of components in the GMMs were varied from two to nine, and the Bayesian information criterion (BIC) was used to select the appropriate model (given the criterion’s ability to penalise model complexity along with assessing the model’s fit). The Bayesian information criterion is,

$$BIC = m \ln N - 2 \ln \hat{L} \tag{21}$$

where m is the number of parameters in the model, N the number of data points used in training and \hat{L} the estimate of the complete data likelihood of the unsupervised GMM. Figure 8 presents the BIC for each of the models considered, where ten repeats were performed to account for the variability in the EM algorithm from the random initialisations of




Fig. 8 Comparison of the BIC for unsupervised GMMs, with the number of components ranging from two to nine. The GMMs were trained on the joint Z24 and KW51 datasets in the transformed space. The bars denote the minimum BIC over ten random initialisations, with the error bar denoting the maximum estimated BIC

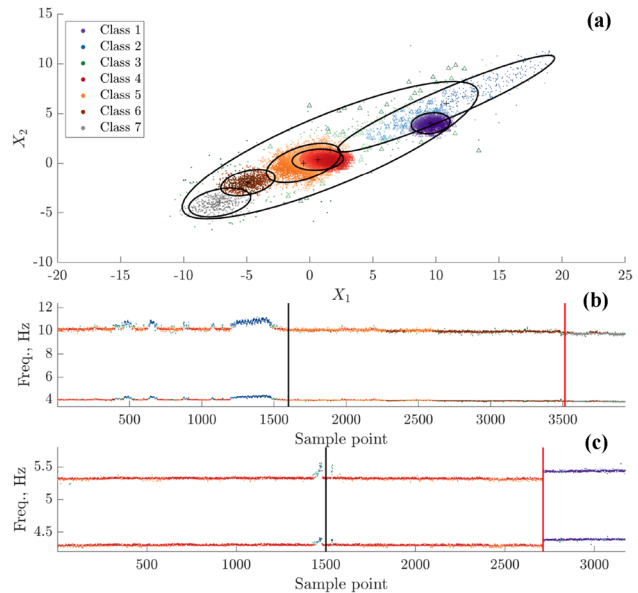


Fig. 9 Unsupervised seven-component Gaussian mixture model predictions on the transformed Z24 and KW51 datasets. Panel (a) is a comparison of the two features for each bridge against the probability of being in one of seven classes, with the inferred (source and target) unsupervised Gaussian mixture model for reference (μ (+) and 2Σ (—)). The Z24 is denoted X_s (·) and the KW51 \hat{X}_t (Δ). Panels (b) and (c) are the Z24 (ω_1 and ω_3) and KW51 (ω_{10} and ω_{12}) natural frequencies against sample point. The black vertical lines in panels (b) and (c) denote the separation of the training data in the DA-GMM mapping, and the red vertical lines indicate the beginning of damage on the Z24 bridge — lowering of the pier by 20mm which occurred on 10/08/1998 — and the start of the retrofit state for the KW51 bridge

the parameters. A seven-component mixture model produced the lowest BIC and, therefore, was selected as the most appropriate model. It is noted that analysis using an infinite Gaussian mixture model (also known as a Dirichlet process

Gaussian mixture model) on the complete four-dimensional Z24 dataset, automatically selected seven distinct Gaussian clusters [30], highlighting that a seven-component model is appropriate for explaining the joint Z24 and KW51 datasets on the aligned two-dimensional feature space. The seven-component model with the lowest BIC out of the ten repeats was used to analyse the bridge datasets.

The predictions from the unsupervised GMM are displayed in Fig. 9. The results indicate that, once a mapping is inferred, both bridges can be diagnosed together, with class information being shared between both bridges. Fig. 9b and c show that the key damage and retrofit states have been well identified, with the GMM identifying the onset of the smallest damage extent on the Z24, a lowering of the pier by 20 mm on the 10th August 1998 (red line). This result is a significant achievement, given that the model has been inferred on a reduced dataset (only the first and third natural frequencies), and provides state-of-the-art performance when compared to existing methods applied to the Z24 in the literature [23–32]. As found in previous analysis on the Z24 dataset using an infinite Gaussian mixture model [30], the Z24 ambient normal condition is non-Gaussian and can be captured by two additional components; classes Five and Six. These Gaussian clusters show that the normal ambient condition (Class Four) starts to rotate and drift as the bridge experiences warmer environmental conditions (see Fig. 4), which are captured by classes Five and Six. In addition, the low temperature condition can be described by more than one Gaussian component [30], classes Two and Three in this model. Finally, it is noted that the feature space inferred by the DA-GMM has retained physical meaning, something that is particularly useful in avoiding problems associated with negative transfer, i.e. the damage class shows a softening affect from the origin, and the retrofit shows a new stiffened behaviour.

6 Conclusions

Population-based SHM is a branch of structural health monitoring that seeks to utilise information from across a population of structures, to improve diagnostic capabilities. Specifically, PBSHM seeks to transfer label information between members of the population. One method of transferring label information is via domain adaptation, where a mapping can be inferred such that the source and target datasets are harmonised, meaning a classifier trained on one domain will generalise to others in the population. This paper developed and demonstrated the potential of a domain adaptation approach constructed from a Gaussian mixture model formulation, for use in an PBSHM context; namely the domain-adapted Gaussian mixture model.

The domain-adapted Gaussian mixture model seeks to identify a linear mapping from a target dataset onto a source Gaussian mixture model. The approach was demonstrated on three datasets. The first was a numerical dataset,

displaying the method's ability to infer linear mappings in the form of two-dimensional rotations. The second dataset presented a numerical case study involving a heterogeneous population of two shear-building structures. The approach was shown to outperform naively applying the source GMM to the untransformed target dataset, and to be an improvement on an unsupervised GMM trained on the target domain. The final dataset, involving the Z24 and KW51 bridges, showed the approach in a completely unsupervised setting. An unsupervised GMM was inferred from the Z24 dataset, which was utilised in identifying a mapping for the KW51 bridge dataset onto the Z24 dataset. The inferred mapping aligned the two datasets allowing class information to be shared between the two bridges. The inferred feature space also retained physical meaning, something not possible with many of the existing domain adaptation technologies. The case study also demonstrated the potential of a PBSHM approach, even when no labels are known initially; as any future labels from one structure can be used in diagnosing the others in the population.

The DA-GMM presents an alternative to existing domain adaptation approaches, using a probabilistic framework and inferring a mapping directly from the target to source domains. Future research should seek to extend the mapping to be nonlinear, whether by a regression viewpoint, via a basis function approach or the kernel trick, or whether via normalising flows [34]. In addition, more robust estimates of the projection parameters and mixing proportions could be obtained if a Bayesian model was constructed by introducing priors on both these parameters. The model could then be solved using a variational inference approach that may be more robust to initial conditions. The DA-GMM is a promising tool for PBSHM and could also be used for more traditional concept drift scenarios on a single structure.

Acknowledgements Dr. P. Gardner, Dr. N. Dervilis and Prof. K. Worden would like to acknowledge the support of the UK Engineering and Physical Sciences Research Council via grants EP/R006768/1, EP/R004900/1 and EP/R003645/1, respectively. Dr. 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. The authors would also like to thank Dr. Kristof Maes, Dr. Edwin Reynders and Prof. Geert Lombaert for helpful conversations in writing this paper and for providing access to the KW51 dataset.

Declarations

Conflict of interest The authors have no competing interests to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes

were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bull LA, Gardner P, Gosliga J, Dervilis N, Papatheou E, Maguire AE, Campos C, Rogers TJ, Cross EJ, Worden K (2021) Foundations of population-based structural health monitoring, part I: homogeneous populations and forms. *Mech Syst Signal Process* 148:107141
- Gosliga J, Gardner P, Bull LA, Dervilis N, Worden K (2021) Foundations of population-based structural health monitoring, part II: heterogeneous populations and structures as graphs, networks, and communities. *Mech Syst Signal Process* 148:107144
- Gardner P, Bull LA, Gosliga J, Dervilis N, Worden K (2021) Foundations of population-based structural health monitoring, part III: heterogeneous populations, transfer and mapping. *Mech Syst Signal Process* 149:107142
- Tsialiamanis G, Mylonas C, Chatzi E, Dervilis N, Wagg DJ, Worden K (2021) Foundations of population-based structural health monitoring, part IV: the geometry of spaces of structures and their feature spaces. *Mech Syst Signal Process* 157:107692
- Pan SJ, Yang Q (2010) A survey on transfer learning. *IEEE Trans Knowl Data Eng* 22:1345–1359
- Weiss K, Khoshgoftaar TM, Wang D (2017) A survey of transfer learning. *J Big Data* 3:29
- Zhuang F, Qi Z, Duan K, Xi D, Zhu Y, Zhu H, Xiong H, He Q (2021) A comprehensive survey on transfer learning. *Proc IEEE* 109(1):43–76
- Zhang W, Peng G, Li C, Chen Y, Zhang Z (2017) A new deep learning model for fault diagnosis with good anti-noise and domain adaptation ability on raw vibration signals. *Sensors* 17(2):425. <https://doi.org/10.3390/s17020425>
- Gu J, Wang Y (2019) A cross domain feature extraction method for bearing fault diagnosis based on balanced distribution adaptation. In: 2019 Prognostics and System Health Management Conference (PHM-Qingdao), pp. 1–5
- Li X, Zhang W, Ding Q, Sun J-Q (2019) Multi-layer domain adaptation method for rolling bearing fault diagnosis. *Signal Process* 157:180–197
- Wang Q, Michau G, Fink O (2019) Domain adaptive transfer learning for fault diagnosis. In: 2019 Prognostics and System Health Management Conference (PHM-Paris), pp. 279–285
- Gardner P, Liu X, Worden K (2020) On the application of domain adaptation in structural health monitoring. *Mech Syst Signal Process* 138:106550
- Fink O, Wang Q, Svensén M, Dersin P, Lee W-J, Ducoffe M (2020) Potential, challenges and future directions for deep learning in prognostics and health management applications. *Eng Appl Artif Intell* 92:103678
- Gardner P, Bull LA, Dervilis N, Worden K (2021) Overcoming the problem of repair in structural health monitoring: metric-informed transfer learning. *J Sound Vib* 510:116245
- Zhang Z, Chen H, Li S, An Z, Wang J (2020) A novel geodesic flow kernel based domain adaptation approach for intelligent fault diagnosis under varying working condition. *Neurocomputing* 376:54–64
- Chakraborty D, Kovvali N, Chakraborty B, Papandreou-Suppapola A, Chattopadhyay A (2011) Structural damage detection with insufficient data using transfer learning techniques. In: *Sensors and smart structures technologies for civil, mechanical, and aerospace systems*, p 798147. <https://doi.org/10.1117/12.882025>
- Ye J, Kobayashi T, Tsuda H, Murakawa M (2017) Robust hammering echo analysis for concrete assessment with transfer learning. In: *Proceedings of the the 11th International Workshop on Structural Health Monitoring*, pp. 943–949
- Paaßen B, Schulz A, Hahne J, Hammer B (2018) Expectation maximization transfer learning and its application for bionic hand prostheses. *Neurocomputing* 298:122–133
- Maeck J, De Roeck G (2003) Description of Z24 benchmark. *Mech Syst Signal Process* 17(1):127–131
- Maes K, Lombaert G (2021) Monitoring railway bridge KW51 before, during, and after retrofitting. *J Bridge Eng* 26(3):04721001
- Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the em algorithm. *J R Stat Soc Series B Methodol* 39(1):1–22
- Christides S, Barr ADS (1984) One-dimensional theory of cracked Bernoulli-Euler beams. *Int J Mech Sci* 26(11–12):639–648
- Maeck J, Peeters B, De Roeck G (2001) Damage identification on the Z24 bridge using vibration monitoring. *Smart Mater Struct* 10(3), 512–517
- Peeters B, De Roeck G (2001) One-year monitoring of the Z24-bridge: environmental effects versus damage events. *Earthq Eng Struct Dyn* 30(2):149–171
- De Roeck G (2003) The state-of-the-art of damage detection by vibration monitoring: the SIMCES experience. *J Struct Control* 10(2):127–134
- Teughels A, De Roeck G (2004) Structural damage identification of the highway bridge Z24 by FE model updating. *J Sound Vib* 278(3):589–610
- Reynders E, De Roeck G (2010) A local flexibility method for vibration-based damage localization and quantification. *J Sound Vib* 329(12), 2367–2383
- Reynders E, Wursten G, De Roeck G (2014) Output-only structural health monitoring in changing environmental conditions by means of nonlinear system identification. *Struct Health Monit* 13(1):82–93
- Langone R, Reynders E, Mehrkanoun S, Suykens JAK (2017) Automated structural health monitoring based on adaptive kernel spectral clustering. *Mech Syst Signal Process* 90:64–78
- Rogers TJ, Worden K, Fuentes R, Dervilis N, Tygesen UT, Cross EJ (2019) A Bayesian non-parametric clustering approach for semi-supervised structural health monitoring. *Mech Syst Signal Process* 119:100–119
- Bull LA, Rogers TJ, Wickramarachchi C, Cross EJ, Worden K, Dervilis N (2019) Probabilistic active learning: an online framework for structural health monitoring. *Mech Syst Signal Process* 134:106294
- Hughes AJ, Bull LA, Gardner P, Barthorpe RJ, Dervilis N, Worden K (2022) On risk-based active learning for structural health monitoring. *Mech Syst Signal Process* 167:108569
- Maes K, Van Meerbeeck L, Reynders EPB, Lombaert G (2022) Validation of vibration-based structural health monitoring on retrofitted railway bridge KW51. *Mech Syst Signal Process* 165:108380
- Rezende D, Mohamed S (2015) Variational inference with normalizing flows. In: *Proceedings of the 32nd International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 37, pp. 1530–1538. PMLR, Lille, France

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.