



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/190069/>

Version: Published Version

Article:

Poole, J., Gardner, P., Dervilis, N. et al. (2023) On statistic alignment for domain adaptation in structural health monitoring. *Structural Health Monitoring*, 22 (3). pp. 1581-1600. ISSN: 1475-9217

<https://doi.org/10.1177/14759217221110441>

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.

On statistic alignment for domain adaptation in structural health monitoring

Jack Poole¹ , Paul Gardner¹, Nikolaos Dervilis¹, Lawrence Bull² and Keith Worden¹

Structural Health Monitoring

2022, Vol. 0(0) 1–20

© The Author(s) 2022



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/14759217221110441

journals.sagepub.com/home/shm

Abstract

The practical application of structural health monitoring is often limited by the availability of labelled data. Transfer learning – specifically in the form of domain adaptation (DA) – gives rise to the possibility of leveraging information from a *population* of physical or numerical structures, by inferring a mapping that aligns the feature spaces. Typical DA methods rely on nonparametric distance metrics, which require sufficient data to perform density estimation. In addition, these methods can be prone to performance degradation under class imbalance. To address these issues, statistic alignment (SA) is discussed, with a demonstration of how these methods can be made robust to class imbalance, including a special case of class imbalance called a partial DA scenario. Statistic alignment is demonstrated to facilitate damage localisation with no target labels in a numerical case study, outperforming other state-of-the-art DA methods. It is then shown to be capable of aligning the feature spaces of a real heterogeneous population, the Z24 and KW51 bridges, with only 220 samples used from the KW51 Bridge. Finally, in scenarios where more complex mappings are required for knowledge transfer, SA is shown to be a vital pre-processing tool, increasing the performance of established DA methods.

Keywords

Domain adaptation, transfer learning, population-based structural health monitoring, damage localisation, machine learning, deep learning

Introduction

The data-driven approach to structural health monitoring (SHM) is often limited by availability of labelled data.¹ Unsupervised machine learning techniques have been used to detect the presence of damage by utilising outlier analysis.² However, SHM-frameworks that allow for data to be categorised as a specific health-state largely rely on supervised statistical models, preventing their application in practical scenarios. This issue motivates the leveraging of information from a population of structures, giving rise to a new discipline – *population-based SHM* (PBSHM).^{3–5} By considering labelled information across a population of structures, it is more likely that the shared label information would be adequate to facilitate more reliable and detailed diagnostics.

A critical challenge in PBSHM is that data acquired from different structures will differ in their underlying distributions, invalidating the assumption that the training and testing data were drawn from the same distribution – an assumption that conventional machine learning models typically make.⁶ Variation in the data distributions will occur for a number of reasons; in *homogeneous* populations (i.e. nominally identical structures), manufacturing

variations and operating conditions will lead to differences. Populations of structures that cannot be considered nominally identical – *heterogeneous* populations⁴ – may also differ significantly because of structural differences. One approach is to utilise information from across a population is to infer a mapping that aligns the feature spaces.

Domain adaptation (DA) – a sub-field of *transfer learning* (TL) – aims to achieve this goal by learning a mapping that aligns the distributions of the feature spaces,⁷ illustrated in [Figure 1](#). As such, a statistical model that will generalise well to a target structure (target domain), can be learnt by using data from a source structure (source domain). Typically, these technologies are based on distribution distance metrics, which learn a mapping via nonlinear kernels^{8–11} or deep neural networks (DNN).^{12–18}

¹Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Sheffield, UK

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

Corresponding author:

Jack Poole, Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK.
Email: jpoole4@sheffield.ac.uk

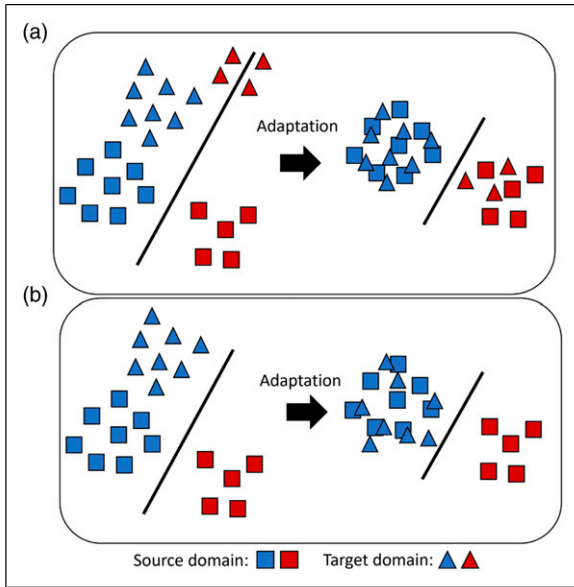


Figure 1. Domain adaptation scenarios showing standard domain adaptation in a and partial domain adaptation in b.

Other approaches align subspaces, assuming that they lie on a common manifold.¹⁹

These technologies have the potential to advance SHM and facilitate more informative diagnostics. Chakraborty *et al.* used TL to address the issue of sensor coverage in a fatigue damage case.²⁰ This application differs from the current paper, which aims to transfer information between structures. A large portion of the literature focusses on DNN based DA, notably Xu *et al.* use a physics-informed deep-DA approach to align the frequency response of multiple structures to perform damage detection and quantification.²¹ There have also been a number of attempts to apply deep-DA to perform fault diagnosis in machines under changing loading conditions and rotation speeds.^{22–27} These works focus on end-to-end DNNs (often convolutional neural networks), with high-dimensional frequency responses used as features; thus, because of the curse of dimensionality⁶ and the large number of parameters in these networks, large quantities of data are required to train these methods such that they generalise well. Conversely, this paper aims to develop an approach to DA that can be applied with limited data, as engineering datasets, particularly SHM datasets, are often sparse and data are expensive to obtain. Another approach that has been investigated focuses on fine-tuning DNNs for image-based crack detection.^{28,29} This approach differs from unsupervised DA, as it does not align the domains and requires target labels.

Domain adaptation has also been demonstrated in a PBSHM setting, Gardner *et al.* have shown that DA can be used to transfer localisation labels between numerical and experimental structures,³⁰ two heterogeneous aircraft

wings³¹ and between pre- and post-repair states in aircraft wings.³² In Bull *et al.*, a population of six experimental tailplanes was used to demonstrate transferring a damage detector.³³ These works have largely focused on methods that rely on kernel-based nonparametric density estimation. Nonparametric density estimation is known to require large quantities of data to provide accurate estimates,³⁴ which may limit these DA approaches in practical SHM scenarios.

In addition, the existing approaches do not address issues relating to class imbalance and partial DA. Class imbalance, where the data quantity for each class is inconsistent, is inevitable in SHM datasets because of the rarity of damage and spurious environmental conditions.³⁵ A specific class imbalance scenario in TL called *partial DA*,³⁶ refers to the scenario where the target domain contains data pertaining to fewer classes than the source, demonstrated in Figure 1. It is crucial that DA methods that are robust to these issues are developed since previously used methods are prone to performance degradation in these scenarios.^{37,36} For fault detection in machines, instance-weighting has been used to achieve impressive classification improvement for partial DA.^{38,39} However, these approaches also rely on DNNs and to the best of the authors knowledge partial DA remains uninvestigated in the context of SHM.

The main idea of this paper is that *statistic alignment* (SA),⁴⁰ a branch of DA that directly aligns the lower-order statistics, can be used as a low-risk form of DA in scenarios with limited data availability and engineering insight can be used to adapt SA to address the issues with class imbalance and partial DA. The main contributions of this paper are as follows:

1. Statistic alignment is shown to perform effective knowledge transfer where standard SHM, and the major branches of DA methods utilised in prior work (manifold, kernel and DNNs) are ineffective when standard practices for standardisation are followed.
2. An approach to SA utilising only normal-condition data that is robust to class imbalance and partial DA scenarios is proposed, along with two methods: normal-condition alignment (NCA) and normal-correlation alignment (NCORAL). These methods are shown capable of adapting both numerical (shear structures) and real heterogeneous populations (the Z24 and KW51 bridges), demonstrating that this approach could facilitate the use of supervised or semi-supervised machine learning methods before any damage has occurred in the target structure.
3. The limitations of SA are discussed; a case study demonstrates that, when the higher-order statistics differ significantly, further DA can be beneficial. It is found that in this scenario, each of the major branches of DA fail to adapt the domains following standard practice for pre-processing, but they

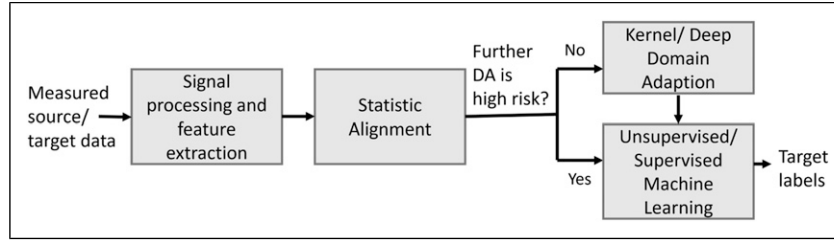


Figure 2. The proposed flow for domain adaptation in PBSHM.

Note: some DA methods may incorporate domain adaptation and a predictive function (i.e. a classifier) in a single model.

succeed following SA, suggesting that SA is an essential pre-processing step for the use of these methods.

The proposed approach to DA in PBSHM is given in Figure 2. The results in this paper suggest that SA should be used as initial adaptation, after which a statistical model can be learnt directly on the source domain, or further, DA could improve adaption prior learning a statistical model. When applying DA, there is a risk that performance is worse than using the target alone;⁴¹ thus, this paper discusses why SA is a low-risk form of DA in low data scenarios and when further DA may be appropriate.

The outline of this paper is as follows. In section 2, the necessary background is given on DA and SA. The disadvantages of conventional DA are discussed in regard to data availability, class imbalance and partial DA, and the proposed approach that aligns the domains using only normal-condition data is introduced. Section 3 demonstrates that SA can transfer label information between numerical structures in both conventional DA and partial DA, where previously investigated DA methods fail, and it is shown that the proposed methods are particularly robust for partial DA. In Section 4, the application of NCORAL is discussed for a real population bridges, the Z24⁴² and KW51 bridges.⁴³ This population presents two partial DA scenarios, which include three domains pertaining to the two heterogeneous bridges, as well as, a pre- and post-repair state in the KW51 bridge dataset. Section 5 discusses the limitations of only aligning the lower-order statistics and introduces SA as a pre-processing step for other prominent DA algorithms, with another numerical case study suggesting that SA may be an essential pre-processing step for the application of many DA algorithms. Finally, Section 6 presents a discussion on SA and discusses future work. A GitHub repository accompanies this paper: <https://github.com/Jack-Poole/Statistic-Alignment.git>.

Domain adaptation

Before defining DA, it is important to introduce two key objects in TL, a domain and task:⁷

- A *domain* $\mathcal{D} = \{\mathcal{X}, P(X)\}$, defined by of a feature space \mathcal{X} and a marginal probability distribution $P(X)$.
- A *task* for a given domain is defined by $\mathcal{T} = \{\mathcal{Y}, f(\cdot)\}$, where \mathcal{Y} is the label space and $f(\cdot)$ is a predictive function learnt from a finite sample $\{\mathbf{x}_i, y_i\}_{i=1}^n$, where $\mathbf{x}_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$.

In *unsupervised DA*, a source domain $\mathcal{D}_s = \{\mathbf{x}_{s,i}, y_{s,i}\}_{i=1}^{n_s}$, with n_s source instances $\mathbf{x}_{s,i}$, each with labels $y_{s,i}$ and a target domain $\mathcal{D}_t = \{\mathbf{x}_{t,j}\}_{j=1}^{n_t}$ with n_t unlabelled target instances, $\mathbf{x}_{t,j}$, are used to learn a classifier that generalises to the target domain. It is assumed that there are differences in the marginal distributions $P(X_s) \neq P(X_t)$, and/or the conditional distributions $P(\mathbf{y}_s|X_s) \neq P(\mathbf{y}_t|X_t)$. Thus, DA aims to find a mapping that aligns the data distributions. The focus of this paper is homogeneous DA, where the feature spaces are assumed the same $\mathcal{X}_s = \mathcal{X}_t$.

This paper is also concerned with *partial DA* problems, where the available target data pertains to fewer classes than the source, that is, the target label space is a subset of the source $\mathcal{Y}_t \subset \mathcal{Y}_s$.⁷ For example, consider a discrete damage localisation problem with five locations, the source may have data for all five locations, whereas the target could only have data pertaining to one location, where the aim of using partial DA would be to use the data in the source domain to learn what location the damage in the target relates to.

A major consideration when attempting unsupervised DA is that the lack of labels in the target domain makes validation challenging. Reducing the risk of *negative transfer*,⁴¹ where DA causes performance degradation, is an ongoing focus of TL research and is of particular importance in SHM where safety critical and high cost assets are at risk.

Statistic alignment

Domain adaptation methods typically attempt to completely harmonise the data distributions, often utilising nonparametric distance metrics to find a nonlinear mapping.⁷ Generally, these metrics rely on the data in both domains being sufficient to perform accurate density estimation. However, in SHM problems, data from a given structure are likely to be limited, particularly for more informative data,

such as damage-states or disparate environmental conditions. A branch of DA, *SA*, provides an alternative solution to align the datasets; and, it is defined as,

Definition 1. *Statistic alignment* methods for DA directly align the lower-order statistics of the domains via affine transformations and scaling.

These approaches typically focus on matching the first- and second-order statistics, and are most appropriate when

- Data can be assumed Gaussian, which is a common assumption in SHM for the linear response of a structure.¹
- The higher-order statistics are already similar.
- Data are limited as robustly estimating the lower-order statistics requires significantly less data than nonparametric density estimation.

In addition to being less data intensive, SA could also facilitate better visualisation of the domain data. Many prominent DA methods^{7,44} project the data into a latent space via a nonlinear mapping. In comparison, SA maintains the structure of the original feature space, as it is limited to affine transformations; this can be useful for physically-interpretable features, which are common in SHM, because features often correspond to some physical process; for example, an increase in natural frequency can be interpreted as a stiffening effect.

One of the most popular SA methods is *correlation alignment* (CORAL),⁴⁵ which aligns the source correlation with the target. The multiple outlook mapping algorithm is a similar approach that aligns the principal components.⁴⁶ He et al. proposed Euclidean alignment (EA), which aligns the means of the covariance matrices for 2D electroencephalogram features, and demonstrated that SA methods can be generalised to cases with multiple sources.⁴⁷

A related branch of DA concerns batch normalisation approaches,⁴⁸⁻⁵⁰ which align the means and standard deviations of the *activations* in DNNs. These methods differ to SA, as the statistics of the original data are not directly aligned and the objective of the network (classification etc.) will influence the feature representation.

Standardisation as SA

Standardisation, a form of normalisation, is commonly used in conventional machine learning to give each feature equal treatment.⁶ In DA, it also has the potential to align the means μ and standard deviations σ of the marginal distributions $P(X_s)$ and $P(X_t)$ in an unsupervised manner, with the following transformations

$$\mathbf{z}_s^{(j)} = \frac{\mathbf{x}_s^{(j)} - \boldsymbol{\mu}_s}{\boldsymbol{\sigma}_s} \quad (1)$$

$$\mathbf{z}_t^{(j)} = \frac{\mathbf{x}_t^{(j)} - \boldsymbol{\mu}_t}{\boldsymbol{\sigma}_t} \quad (2)$$

where $\boldsymbol{\mu}_s$ and $\boldsymbol{\mu}_t$ are the means of the source and target; $\boldsymbol{\sigma}_s$ and $\boldsymbol{\sigma}_t$ are the respective standard deviations. In this paper, this form of standardisation will be referred to as *A-standardisation*. On the other hand, standard practices for machine learning would suggest that the statistics applied to all data should be the same; for example, if some data from the source and target domain are considered training data the statistics would be calculated from $X = X_s \cup X_t$, which may remove the effect of measurement scale without changing relative mean distance or scale between the domains. This method will be called *N-standardisation* in this paper.

Previous work has studied various approaches to normalisation to improve the performance of a popular DA algorithm, transfer component analysis (TCA),⁵¹ showing standardisation can lead to better adaptation. However, it suggests a range of approaches, including methods that would not align the statistics (such as N-standardisation), and does not consider using standardisation alone to perform adaptation. The current paper first considers standardisation as a form of DA itself, before considering it as a pre-processing method.

Correlation alignment

A-standardisation aligns the domains, ignoring the correlation between features. Correlation Alignment,⁴⁵ extends this method to also align the covariance. This is achieved by transforming the source domain via a linear transformation matrix A , such that

$$\min_A \|\widehat{C}_s - C_t\|_F^2 = \min_A \|A^T C_s A - C_t\|_F^2 \quad (3)$$

where \widehat{C}_s is the covariance of the transformed source, C_s and C_t denote the covariance matrices of the source and target, respectively, and $\|\cdot\|_F$ is the Frobenius norm. In many SHM cases, CORAL may improve generalisation of the source classifier, but estimating covariance suffers from the curse of dimensionality. If the number of observations in X_s or X_t are smaller than the number of dimensions d ($n < (d + 1)$), the covariance matrix will be singular.⁵² This can be a serious problem since vibration data are often high-dimensional. In damage detection, the curse of dimensionality has motivated the use of ensemble methods to robustly estimate covariance.^{52,53}

Normal-condition alignment

Structural health monitoring datasets are likely to have class imbalance; thus, current SA methods may not be robust to SHM problems. For example, given two bridge datasets it is unlikely that both bridges will have the same quantity of

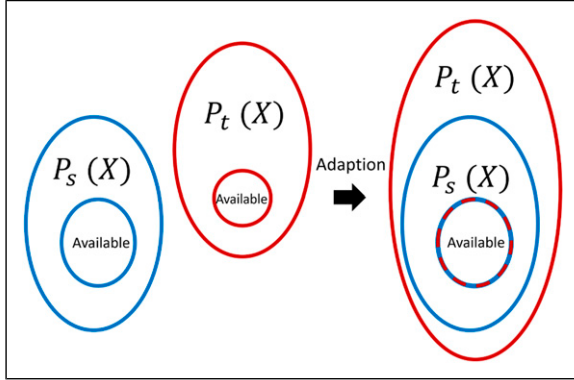


Figure 3. Demonstration of aligning all the available data in the context of partial DA or with class imbalance. The source and target, shown in blue and red, respectively, have available data pertaining to a subset of the underlying distribution, shown by the inner circle. After alignment, the underlying distributions are not well aligned because the target data is representative of a smaller subset of the underlying distribution.

available data from every damage-state and environmental condition. It is also unrealistic to assume the target will contain data from each health-state in the source – motivating the application of partial DA. As such, their statistics summarise different behaviours and aligning the domains based on these moments will not align the underlying distributions, leading to negative transfer, illustrated in Figure 3. Similarly, prominent DA methods typically only minimise the marginal data distributions,⁷ which may also result in a subset of classes in the target being aligned to match the distribution of the entire source. Thus, NCA is proposed, which aims to reduce the risk of aligning data generated via different processes, by utilising the assumption that data gathered at the start of a structures operation were generated by the ‘normal-condition’ – a common assumption made for novelty detection.² In NCA, the source domain is first standardised via equation (1) to centre the data and give features equal treatment. The normal-condition of the target domain is then aligned with that of the source by

$$z_t^{(j)} = \left(\frac{x_t^{(j)} - \mu_{t,n}}{\sigma_{t,n}} \right) \sigma_{s,n} + \mu_{s,n} \quad (4)$$

where $\mu_{s,n}$, $\mu_{t,n}$ and $\sigma_{s,n}$, $\sigma_{t,n}$ are the means and standard deviations of the normal-condition data for the source and target, respectively. If variation between the datasets is assumed to be limited to scale vector \mathbf{a} , and translation vector \mathbf{b} , the differences can be expressed by

$$X_s = \mathbf{a}X_t + \mathbf{b} \quad (5)$$

Given each dataset can be expressed as $X = X_n \cup X_d$ where X_n , is normal-condition data and X_d is damage-state

data and the \mathbf{a} and \mathbf{b} are affine transformations, it follows that the transformations for the entire domain can be learnt from the subsets X_n . It is noted that using only a subset of the data reduces the available data to learn the statistics, but the lower-order statistics should be able to be estimated with a limited sample size.

The main aim of DA is to adapt the underlying distributions so that a given health-state c from each domain follows the same conditional distribution $P_s(Y = c|X) = P_t(Y = c|X)$. Given that the data distributions for a finite sample will contain biases because of class imbalance and differences in the label spaces, this task is challenging. Previous SA and typical DA approaches would naively attempt to align the marginal distributions, which would not correctly align the underlying distributions, shown in Figure 3. Therefore, to align the underlying distributions, a subset of the source label space must be chosen, and the quantity of data in each class should be balanced. Explicitly aligning the marginal distributions of data which are believed to have been generated by the normal-condition is a low-risk way of selecting data corresponding to the same label space. This strategy is in contrast to other approaches, such as automatic weighting or sample selection procedures based on domain discriminators, which may unpredictably select data from different health-states.²⁵

Furthermore, reweighting classes to address class imbalance is challenging, as labels are not available in the target. However, if the data are limited to normal-condition data, it can be assumed that variations are related to measurable environmental and operational conditions (EoVs), such as atmospheric temperature, traffic loading or wind speed etc.⁵⁴ These measurements could allow for further sample selection such that the normal-conditions of each domain contain data from more similar EoVs.¹

The advantages of aligning the domains based on the normal-condition are demonstrated in Figure 4. Figure 4(a) presents the original toy problem, consisting of a source domain with three Gaussian clusters and a target with two classes. Hence, the problem is a partial DA scenario and there are also differences in the class imbalance between the available classes, with 20 samples in Class 0 (shown in red) for both domains, but with eight and four samples in source and target, respectively, for Class 1 (shown in blue). Figure 4(b) presents N-standardisation, showing that it provides no adaptation. In Figure 4(c), A-standardisation has misaligned the standard deviation and mean of the two classes in the target to the three in the source. Normal-condition alignment attempts to address this issue by only considering data from the normal-condition, which in this case aligns the correct classes because the red and blue classes have a similar structure in both domains, shown in Figure 4(d).

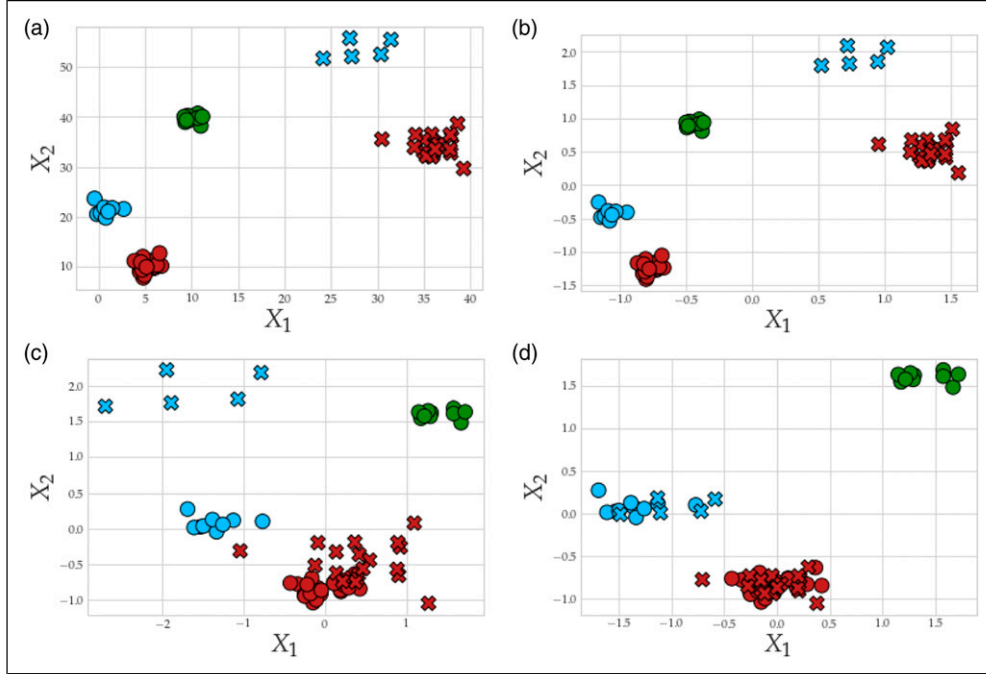


Figure 4. Demonstration of aligning a toy example, panel a, presenting a partial DA problem, consisting of three Gaussian clusters in the source and two in the target. Panel b gives the result of N-standardisation, c gives A-standardisation and d gives NCA. The source and target data are represented by (O) and (x), respectively; Classes 0, 1 and 2 are depicted in red, blue and green, respectively.

Normal-correlation alignment

Correlation alignment may also be prone to negative transfer under class imbalance, as accurately estimating the global correlation would be challenging. A modification of CORAL could exploit information in the correlation between the normal data – NCORAL. The first step of NCORAL is to apply NCA; NCORAL is then given by

$$\min_A \|\widehat{C}_{s,n} - C_{t,n}\|_F^2 = \min_A \|A^T C_{s,n} A - C_{t,n}\|_F^2 \quad (6)$$

where $C_{s,n}$ and $C_{t,n}$ are the correlations of the normal-condition in the source and target, respectively. Normal-Correlation Alignment extends the assumption that the domains differ by a scale and translation made by NCA, by also considering a rotation.

It is noted that NCORAL learns the correlation from a subset of the entire data, so may have additional issues relating to the curse of dimensionality; this issue may be addressed by ensemble methods for high-dimensional features.²

Case study: Numerical three-storey population

In SHM, the application of supervised and semi-supervised models is challenging in practical scenarios because label information is often sparse and expensive to obtain.¹ This

section presents a numerical case study, demonstrating the ability of SA to transfer label information to facilitate damage localisation with a limited quantity of data and no labels in the target domain. A number of SA methods are applied – A-standardisation, CORAL, NCA and NCORAL – and these are benchmarked against, N-standardisation (showing the result of applying traditional SHM methods to a population) and a range of DA methods that encompass the general DA approaches used in prior work – transfer component analysis (TCA),⁸ balanced distribution adaptation (BDA),¹¹ the geodesic flow kernel (GFK)¹⁹ and the domain adversarial neural network (DANN).¹³

Simulation

The population consists of two shear structures modelled as 3DoF lumped – mass models (calculated following the approach in Ref. [30]). The masses of each DoF were assumed to be a rectangular volume, representing a floor, parameterised by a length l_m , width w_m , thickness t_m and density ρ , with the density sampled from a Gaussian distribution to represent manufacturing variation. The masses were assumed connected by four cantilever beams in parallel, so stiffness is given by $k = 4k_b$, where stiffness of each beam was found as the tip stiffness of a cantilever beam, $k_b = \frac{3EI}{l_b^3}$. The elastic modulus E was also drawn

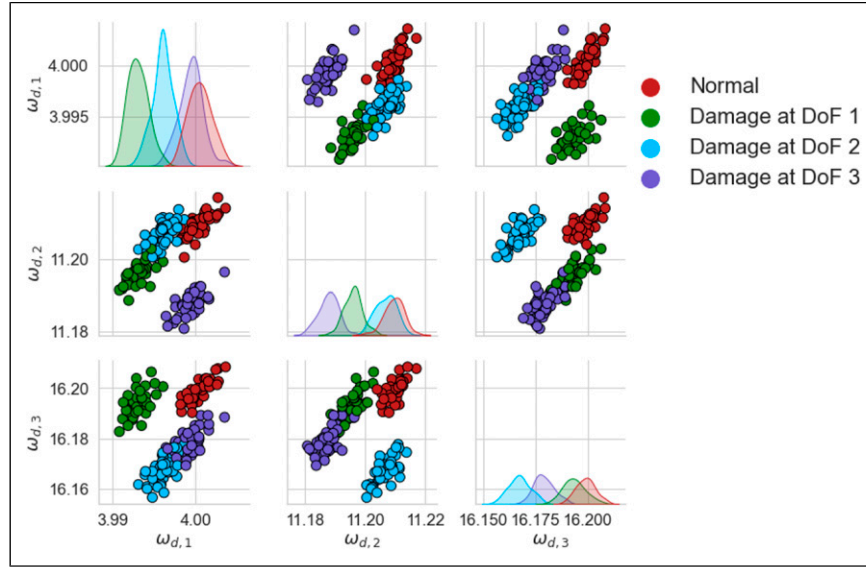


Figure 5. Unnormalised damped natural frequencies of the source structure of the numerical population of three-storey, in Hz. A random subset 20% the size of the dataset was used for visualisation.

from a Gaussian distribution for each sample to introduce variability. Damping c was not derived from a physical model; instead, it was drawn from a Gamma distribution directly.

Damage at a given storey was modelled as an open crack on one of the four beams, located at the midpoint of the beam. It was modelled as a reduction in stiffness as in Ref. [55]; thus, $k = k_d + 3k_b$, where k_d is the stiffness of a damaged cantilever beam.

Having obtained the parameters of the model, the damped natural frequencies ω_d were calculated by solving the eigenvalue problem; the first three were used as features, $X \in \mathbb{R}^{n \times 3}$.

Material properties and geometry for each structure are detailed in the Appendix. Data were simulated for the normal-condition and three damage classes, representing damage located at each DoF. For the source structure, 200 samples were collected for each class and labels were assumed known $\{\mathbf{x}_{s,i}, y_{s,i}\}_{i=1}^{n_s}$, where $n_s = 800$. In the target structure, 100 samples were collected for each class $\{\mathbf{x}_{t,j}\}_{j=1}^{n_t}$, where $n_t = 400$. The target labels were assumed to be unknown for all classes apart from the normal-condition. In addition, a separate test target dataset was generated via the same procedure as the training set, that is, $\{\mathbf{x}_{test,j}, y_{test,j}\}_{j=1}^{n_{test}}$, where $n_{test} = 400$.

In this paper, the case studies focus on vibration-based SHM, but other features could be utilised. It is interesting to note that previous work has applied DA technologies to the similar field of non-destructive testing with ultrasound data,^{56,57} and both studies find normalisation or batch normalisation that adapts the domains has a key role in achieving transfer.

Comparison

To demonstrate that SA can robustly transfer label information between structures, the labels for the three discrete damage locations, corresponding to each DoF, and normal-condition in the source were transferred to the target to facilitate damage localisation for the same health-states, without using target labels. Four SA methods were considered; alignment via A-Standardisation given by equations (1) and (2), CORAL, NCA and NCORAL. In addition, to benchmark these methods against traditional SHM (no DA), N-Standardisation (i.e. calculating the statistics from $X = X_s \cup X_t$) was applied. After alignment, a k-nearest neighbours classifier (KNN), with one neighbour, was learnt on a source training set and used to classify data in a target test set, although any appropriate classifier can be applied following SA, a KNN is used here, because if the source and target distributions are well aligned, data should be close in Euclidean distance. Furthermore, SA is compared to some prominent DA approaches, using N-standardisation for pre-processing; these are as follows:

- *Transfer component analysis*:⁸ learns a feature space that minimises a nonparametric distribution distance metric, the maximum mean discrepancy (MMD),⁵⁸ between the marginal distributions $P(X_s)$ and $P(X_t)$, by MMD-regularised kernel principal component analysis (PCA).
- *Balanced distribution adaptation*:¹¹ extends TCA by also attempting to minimise the MMD between the conditional distributions $P(\mathbf{y}_s | X_s)$ and $P(\hat{\mathbf{y}}_t | X_t)$, where $\hat{\mathbf{y}}_t$ are label predictions. Since there are no labels in the target, BDA uses pseudo labels from predictions to

assign target samples to a given class and finds the distance between the class-conditional distributions $P(X_s|y_s)$ and $P(X_t|\hat{y}_t)$.

- The *geodesic flow kernel*:¹⁹ uses the ‘kernel trick’ to find a projection of infinite subspaces on the Grassmann manifold that bridge the source and target PCA subspaces, which are assumed to lie on a common manifold.
- The *domain adversarial neural network*:¹³ uses a deep feature extractor to find a domain-invariant space. Taking inspiration from the generative adversarial network,⁵⁹ a domain discriminator is trained to distinguish between the source and target. Gradient ascent is performed in the feature extractor to align the distributions; a classifier is simultaneously trained to ensure the feature space is discriminative.

These methods include the main forms of feature extraction previously used for DA in SHM – linear subspaces, kernels and DNNs – as well as methods that minimise a nonparametric statistical distance metric, a domain discrimination loss (adversarial approach) and the distance along the geodesic.

Hyperparameter selection via cross-validation is challenging in unsupervised DA, because labels in the target are assumed unavailable. As such, the MMD-based methods (TCA and BDA), utilise an RBF kernel with the length scale chosen as the median of the pairwise distances;⁵⁸ the dimension of the feature space is reduced by one and the Frobenius-norm regularisation parameter was arbitrarily chosen as 0.1; results are largely insensitive to this value.⁸ Balanced distribution adaptation includes a ‘balance factor’ to control the contribution of the MMD between the marginal and conditional distributions; this was chosen to be 0.5 as suggested in Ref. [11] for the unsupervised setting. The dimension of the GFK subspace must be 1, since there is a requirement that it is less than half the original dimension.¹⁹ The classifier learnt on the features found by TCA, BDA and the GFK methods was a KNN, with one neighbour. The architecture of the DANN was chosen from a similar case, given in Ref. [60]. The DANN is sensitive to the random initial weights, so 100 repeats were run and the mean and one standard deviation are given.

Results are given for a test dataset in the target domain, where the evaluation metric used is the macro-F1 score (see Ref. [30] for more details).

Results: Numerical three-storey population

The unnormalised natural frequencies of the source and target can be found in Figures 5 and 6, respectively. The differences between the domains are large, with the

absolute values in the target being about a factor of two larger than the source. Estimations of the class data distributions are given by kernel density estimation (see Ref. [61] for more details), shown on the diagonal of each figure.

Results comparing each DA method can be found in Figure 7. As expected, the KNN trained on the N-standardised (unadapted) features, which can be considered as naively applying a classifier trained via a traditional SHM approach to another structure, lead to poor generalisation of the source classifier. Following any SA, perfect classification could be achieved. The features given by A-standardisation can be found in Figure 8. It can be seen that even though the two structures have significant differences, the domains have been well aligned, perhaps suggesting differences between similar structures caused by size and material properties may only lead to differences in mean and scale, assuming a linear response. Furthermore, the conventional DA algorithms did not improve classification performance upon N-standardisation. These methods should be able to align both the lower- and higher-order statistics, so this result may suggest that large differences in scale and mean may make learning challenging in conventional DA algorithms. This issue may be related to the quantity of data. It is noted that as the mean and standard deviation are relatively simple to calculate from the data directly, SA could be used to reduce the mean and scale discrepancies before these algorithms are applied, motivating the idea of SA as a pre-processing tool.

Results: Partial domain adaption with the numerical three-storey population

In practical scenarios for PBSHM, it is likely that the amount of data from each health-state and the labels spaces would differ between domains; in such cases, conventional DA is prone to negative transfer.³⁶ Thus, the use of partial DA methods should be considered. Normal-condition alignment and NCORAL intrinsically addresses the partial DA problem by only adapting the normal-condition. To investigate the robustness of previously applied methods to partial DA and class imbalance, the target domain was downsampled to only include 10 samples from one damage-state, corresponding to damage on the third storey – this is a partial DA problem as the target only has two classes, which are a subset of the four in the source. Therefore, the transfer problem in this case is to transfer the label information from the source structure to the target, given that the target only contains a small quantity of data from one of the three discrete damage locations in the source.

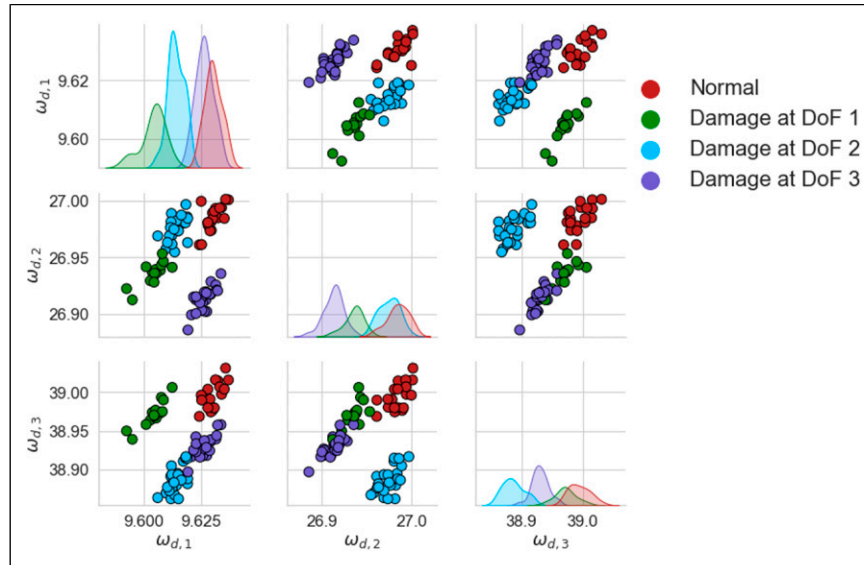


Figure 6. Unnormalised damped natural frequencies of the target structure of the numerical population of three-storey, in Hz. A random subset 20% the size of the dataset was used for visualisation.

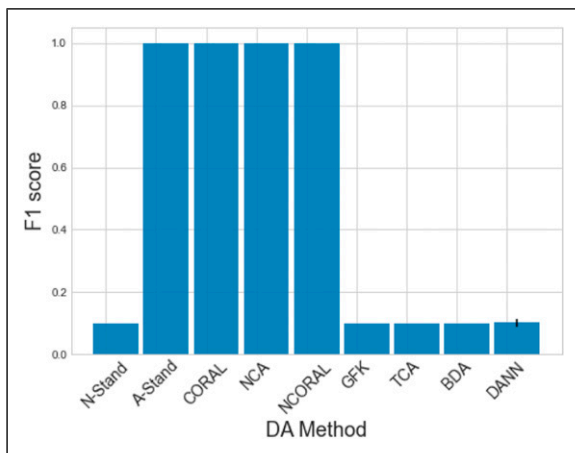


Figure 7. Classification performance of a KNN on the target domain after DA on the numerical three-storey population. The result of the DANNN is given as the mean of 100 repeats with one standard deviation shown by a black line. DANNN: domain adversarial neural network; KNN: k-nearest neighbours classifier.

The results for this case are given in Figure 9. As with the previous case, all the conventional DA methods failed to improve generalisation.² In addition, A-standardisation and CORAL caused negative transfer. Figure 10 shows that aligning the global statistics using A-standardisation in this scenario has caused the two available target classes to be spread across the four classes in the source. Figure 11 shows that by only aligning the normal-condition statistics, the damage data were also aligned. This result could be achieved without any available damage-state data in the

target, allowing damage diagnostics in real time using contextual information from a source domain.

Case Study: Partial DA with the Z24 and KW51 bridges

In a population of real structures, it is more realistic to assume a partial DA setting. This section presents a real population of structure consisting of the Z24,⁴² and KW51 bridges.⁴³ This population consists of two partial DA problems that are challenging to solve using previously investigated DA approaches. First, the KW51 bridge was repaired, which changed the response, so to reuse the previous monitoring data DA must be used to realign the response with the pre-repair data. Second, the Z24 bridge dataset contains damage-state data that could be used to further inform a monitoring system if it can be aligned with the KW51. As such, this section demonstrates how NCA and NCORAL can be used to bring these domains into alignment.

Datasets

The Z24 bridge dataset is well-studied, with data-based approaches being able to identify key events during the monitoring campaign.^{2,62-69} The Z24 bridge was a concrete highway bridge in Bern, Switzerland, which, as part of the SMICES project, was used for an SHM campaign before its demolition in 1998.⁴² The first four natural frequencies were found via operational modal analysis (OMA), from the collected acceleration responses. Small-scale damage was introduced by lowering the pier incrementally on the 10th of

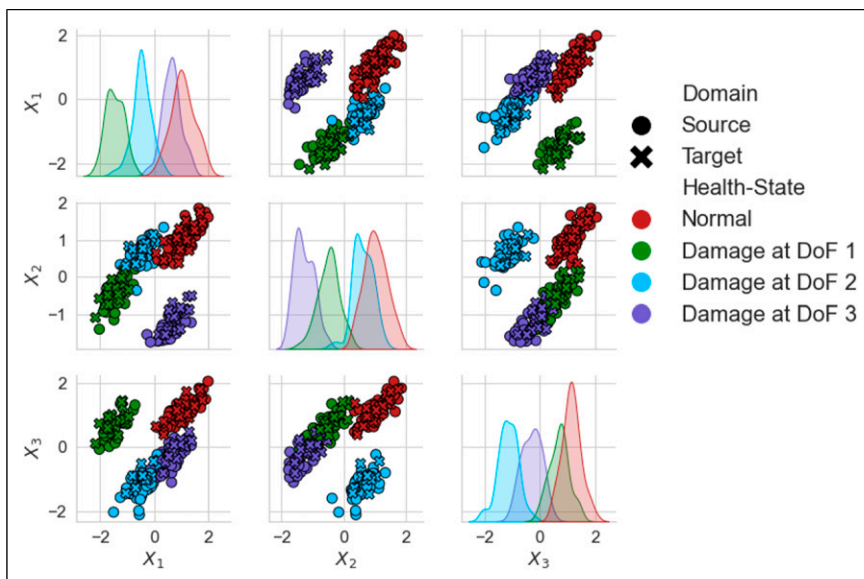


Figure 8. A-standardised features of the numerical three-storey population of structures. The source and target are depicted by (○) and (×), respectively. A random subset 20% the size of the dataset was used for visualisation.

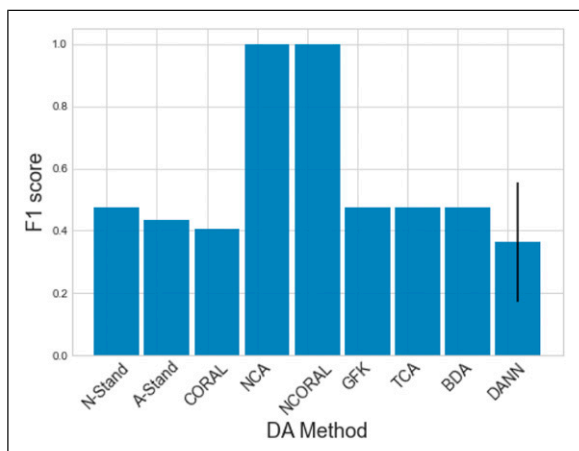


Figure 9. Classification performance of a KNN on the target domain after DA on the numerical three-storey population in a partial DA scenario. The result of the DANN is given as the mean of 100 repeats with one standard deviation shown by a black line. DANN: domain adversarial neural network; KNN: k-nearest neighbours classifier.

August 1998, before more severe damage occurred, starting with the failure of the concrete hinge on the 31st of August 1998. For a more in-depth overview of this dataset, see Ref. [62].

The KW51 bridge is a steel bow-string railway bridge in Leuven, Belgium. A monitoring campaign was carried out between 2018 and 2019 during 15 months. The acceleration responses were used to obtain the first 14 natural frequencies via OMA. During the monitoring campaign, each diagonal member was retrofitted with a steel box to strengthen the

design of the bridge, with the retrofit beginning on the 15th of May 2019 and completed on the 27th of September 2019. Novelty detection of the retrofit has been successfully demonstrated using robust PCA and linear regression trained on the pre-retrofit data.⁷⁰ For a full description of the dataset, the reader is referred to Ref. [43].

Even though the two bridges differ significantly by design, a subset of the natural frequencies can be chosen where there are similarities in the modal response of the bridges. Specifically, the first and third natural frequencies of the Z24 bridge and the 10th and 12th natural frequencies of the KW51 bridge correspond to vertical bending modes of the deck and have to close the same nodal pattern (for visualisation of the mode shapes see, Refs. [42] and [43] for Z24 bridge and KW51 bridge, respectively). The corresponding natural frequencies are visualised in Figures 12 and 13; it can be seen that both bridges experience stiffening effects because of below freezing conditions and the Z24 bridge dataset contains additional information corresponding to damage. Therefore, both bridges can be assumed to have two normal-conditions (normal ambient and low temperature). In addition, it is clear that the KW51 is stiffened by repair, shown by the increase the frequencies for similar temperatures, meaning that the pre- and post-repair states correspond to two domains.

Within the two datasets, there are three domains. Clearly, the response between the Z24 and KW51 bridges are different and require some form of DA for knowledge transfer. In addition, following repair the response of the KW51 bridge changed. This phenomenon has previously been investigated; it was found for pre-repair data to be used to predict the health-state for the post-repair structure, it must

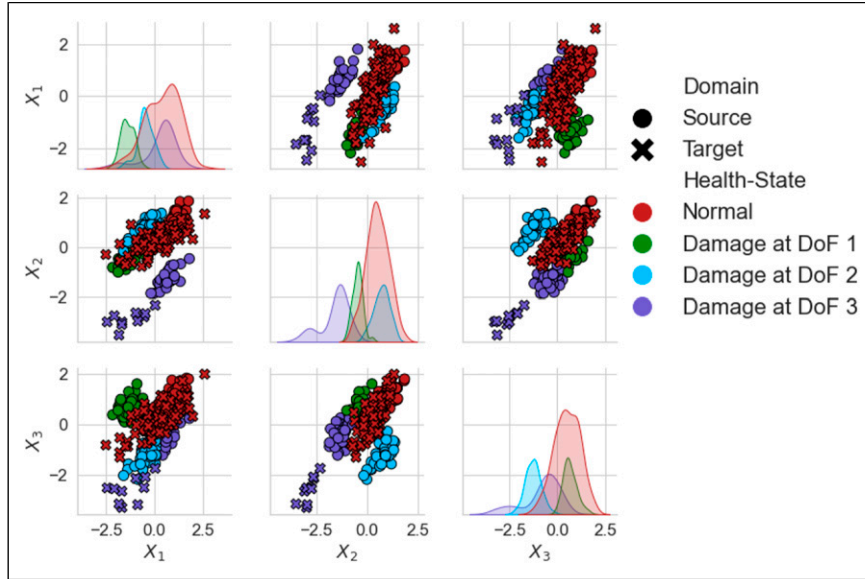


Figure 10. A-standardised features of the numerical three-storey population of structures in a partial DA scenario, with data only from a subset of the source classes in the target. The source and target are depicted by (O) and (x), respectively. A random subset 20% the size of the dataset was used for visualisation.

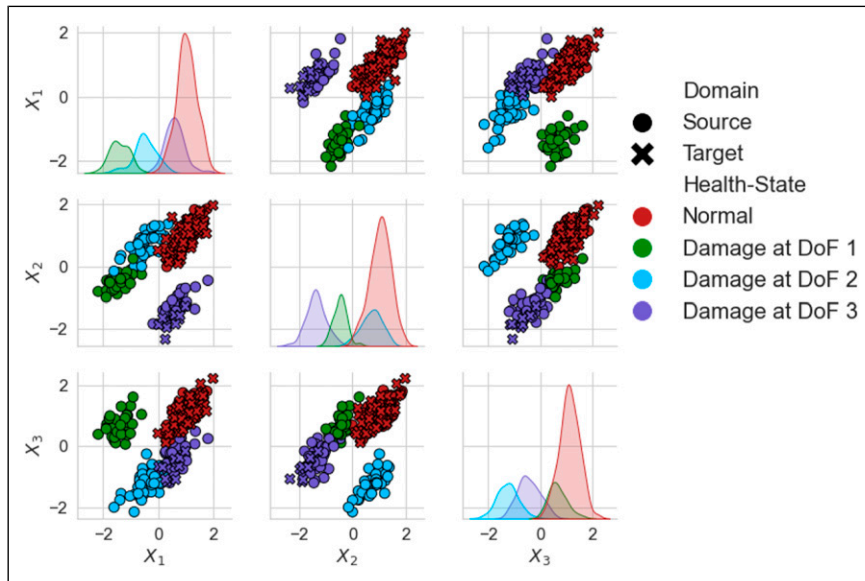


Figure 11. NCA features of the numerical three-storey population of structures in a partial DA scenario, with data only from a subset of the source classes in the target. The source and target are depicted by (O) and (x), respectively. A random subset 20% the size of the dataset was used for visualisation. NCA: normal-condition alignment.

be realigned via DA.³² Thus, to perform future predictions on the KW51 bridge post-repair using the Z24 and KW51 bridge pre-repair data two partial DA problems must be addressed:

1. Align the pre- and post-repair data so that the KW51 bridge data can be treated as one domain. This

problem is defined as one of partial DA, because there are data from the ambient and low temperature normal-conditions in the pre-repair state, but only the ambient normal-condition in the post-repair state.

2. Align the KW51 and Z24 bridge data. This problem is one of partial DA, because there is an additional

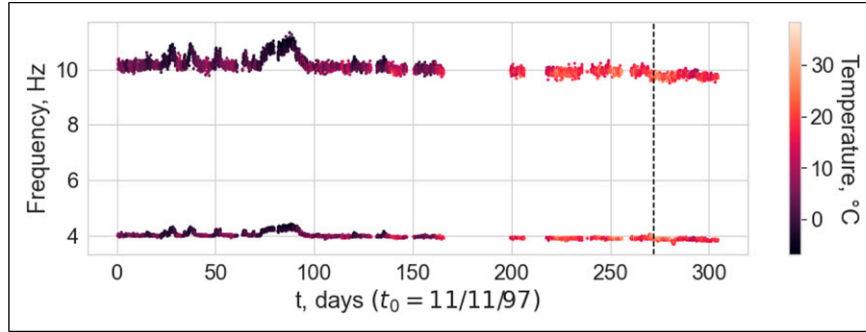


Figure 12. The first (bottom) and third (top) natural frequencies of the Z24 bridge dataset. The first instance of damage, commencing on the 10th of August, is indicated by the black line.

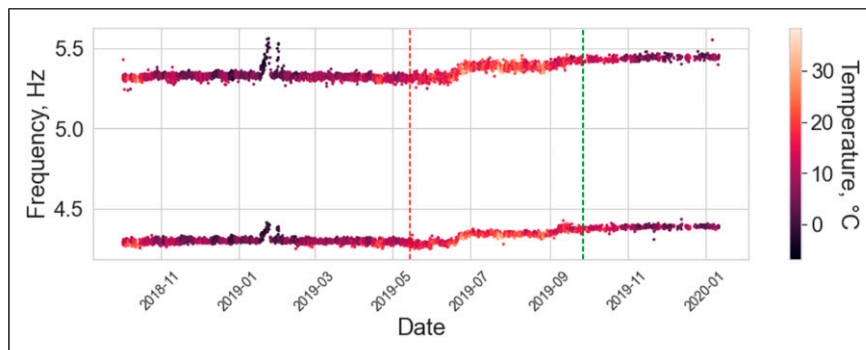


Figure 13. The 10th (bottom) and 12th (top) natural frequencies of the KW51 bridge dataset. The red vertical line indicates the start of the retrofit and the green is the end.

class in the Z24 bridge dataset relating to damage on the deck.

Using previously investigated DA algorithms, this problem would be challenging to solve. The methods that find a subspace (TCA, BDA, the GFK etc.) will reduce the dimension of the original space to unity when adapting the pre- and post-repair domains, and then they could not be used to adapt the KW51 to the Z24 bridges, as the dimension cannot be further reduced.³ In addition, the DNN-based methods often require labels for a classification task in the source. This classification task maintains the discriminative information in the domains. Although, temperature data are available for each bridge, there are no ground-truth labels that indicate the stiffening effects caused by low temperatures. Therefore, there are only noisy labels to maintain discriminative information between the ambient and low temperature normal conditions. Even if these labels are used, there are only 110 samples corresponding to low temperatures in the KW51 bridge dataset, so if an unbiased subset of data were chosen there would be 220 samples in the pre-repair source domain, which is likely too small to train a DNN. As such, this section demonstrates that NCA and NCORAL can be used

to align the pre- and post-repair states of the KW51 before NCORAL is used to align all the data of the KW51 to the Z24 bridge datasets.

Domain adaptation and clustering

Initially, a sensitivity analysis was conducted to evaluate the quantity of data required in each domain for SA. The mean and standard deviations were calculated for each structure with varying sample size, starting at 10 increasing to 500 samples, with a 10 sample step size. The results of the sensitivity analysis are given in Figure 14. It can be seen that the mean can be accurately estimated with very limited data, suggesting transfer could be possible between real structures if the differences are mostly summarised by the mean. As expected, the standard deviations required more data to be accurately estimated, particularly for the KW51.

The first step was to adapt the pre- and post-repair KW51 data. Thus, NCORAL was applied, learning the statistics using the first 200 data from the pre- and post-repair state, which correspond to similar temperatures. Given the sensitivity analysis (Figure 14(d)), this sample size should correspond to a robust estimation of the statistics whilst only

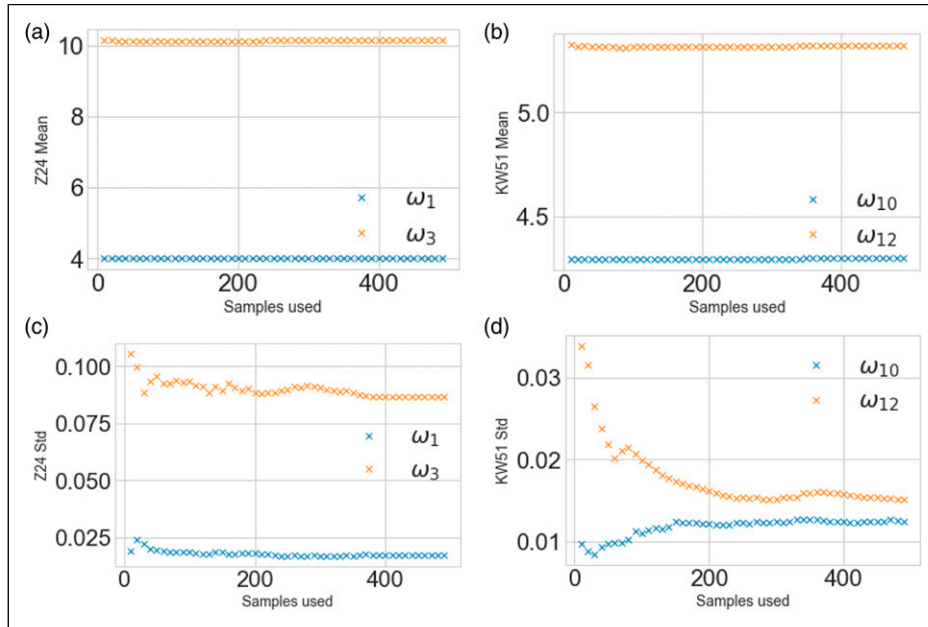


Figure 14. Sensitivity analysis for calculating the means and standard deviations for the Z24 and KW51 bridges for a varying sample size. Panels a and b present the means for the Z24 and KW51 bridges; panels c and d give the standard deviation.

selecting data from a short period after inspection to increase the likelihood that it was generated by the normal-condition.

Once adapted, it was assumed that the pre- and post-repair data could be treated as a single domain. Thus, to align the Z24 and KW51 bridge datasets, NCORAL was applied. As can be seen in Figures 12 and 13, both bridges experience stiffening effects because of freezing temperatures, so a prior assumption is that, within both datasets, normal-condition data can be split into two classes, pertaining to ambient ($T > 0^\circ\text{C}$) and low temperatures ($T < 0^\circ\text{C}$). Labels for these effects are not explicitly known, but temperature data gives an indication as to which data correspond to these effects. As such, data where $T > 0^\circ\text{C}$ was considered ‘ambient normal-condition’ and $T < 0^\circ\text{C}$ as ‘low temperature normal-condition’. Since there is significantly more ambient normal-condition data in both domains, the temperature measurements were used for sample selection. For the Z24 bridge, based on the sensitivity analysis (Figure 14(c)), 400 samples were selected from each normal-condition class, with the low temperature normal-condition being randomly selected from the subset where $T < 0^\circ\text{C}$. In the KW51 dataset, there are only 110 samples corresponding to $T < 0^\circ\text{C}$; so that the dataset is unbiased 110 samples are used from each normal-condition. Based on the sensitivity analysis (Figure 14), this sample size only underestimates the 10th natural frequency by approximately 4% and overestimates the 12th natural frequency by 7%. The KW51 bridge is adapted by only considering a subset of data from the first 72 days of its

monitoring campaign, which can safely to be assumed to be operating in its normal-condition.

To demonstrate that after alignment information between the bridges can be shared, an unsupervised GMM is learnt on the aligned features. Here an unsupervised model is utilised, since ground-truth labels are unknown, but any appropriate model could be applied after SA. The prior assumption is that there are three classes within datasets, the ambient and low temperature normal-conditions and, the Z24 bridge damage, so a three-component model was implemented. Since, ambient normal data are more abundant, to reflect the prior assumption that normal-condition data is split between ambient and low temperature conditions, the temperature data were used to downsample data corresponding to ambient conditions ($T > 0^\circ\text{C}$).⁴

The aligned features and the predictions made by the GMM can be found in Figure 15. It can be seen in Figures 15(b) and (c), that the ambient normal-conditions of the Z24 bridge, pre- and post-repair KW51 bridge are well aligned, as well as, the low temperature normal-condition of the Z24 bridge and pre-repair KW51 bridge. The feature space has maintained physical interpretability (Figure 15(a)), an aspect of SA that could be useful for mitigating the risk of negative transfer. The stiffening effect caused by low temperatures engenders an increase in each feature (in blue, where $T < 0^\circ\text{C}$). In addition, the stiffening reduction caused by damage in the Z24 bridge can be seen by a reduction in each feature (in purple).

One of the main advantages of aligning the population of bridges is that damage-state data from a source could be

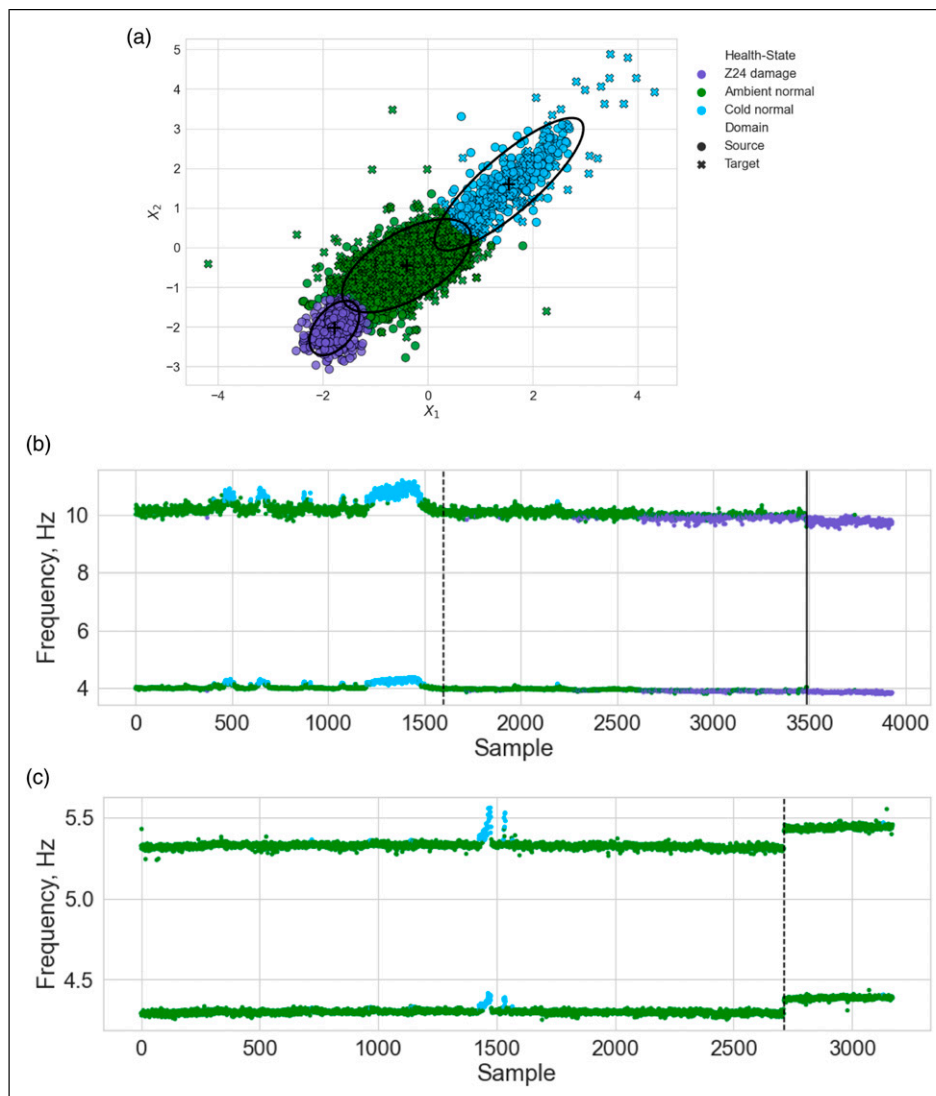


Figure 15. Unsupervised GMM predictions on the Z24 and KW51 bridge datasets. Panel a gives the comparison of the two features after alignment using NCORAL, showing the four Gaussian components identified (μ (+) and 2σ (—)); the Z24 bridge samples are denoted by (○) and the KW51 by (×). Panels b and c gives the predicted classes on the unadapted Z24 bridge (ω_1 and ω_3) and KW51 bridge (ω_{10} and ω_{12}) natural frequencies against sample point.

used to further inform a damage detector for the target. In Figure 15, it can be seen that the normal-condition data for the KW51 lies on the boundary with damage in the Z24 bridge dataset, but no data are misclassified as showing damage.⁵ This result motivates the idea of using damage to inform a novelty detector to reduce the number of false positives that occur when the threshold is defined based on the normal-condition of the target alone. Unfortunately, there is no damage in the KW51 to confirm this finding.

This case study also illustrates a trade-off between selecting features that are transferable and damage sensitive, shown by Figure 15(b) where some normal-

condition data of the Z24 bridge dataset are clustered with damage. It can be seen in Figure 15(a) that this is largely because damage is masked by ambient Z24 bridge data. In previous studies, however, it has been demonstrated that using all four available frequencies allows for damage to be discriminated,² but the most damage sensitive natural frequency in the Z24 bridge dataset (the second natural frequency) was not used since the mode shapes indicate there is low physical similarity with any of the modes in the KW51 bridge. This trade-off may be influenced by the similarity of the population. In this population, a concrete box-girder highway bridge is used to transfer information to a steel bow-string railway bridge,

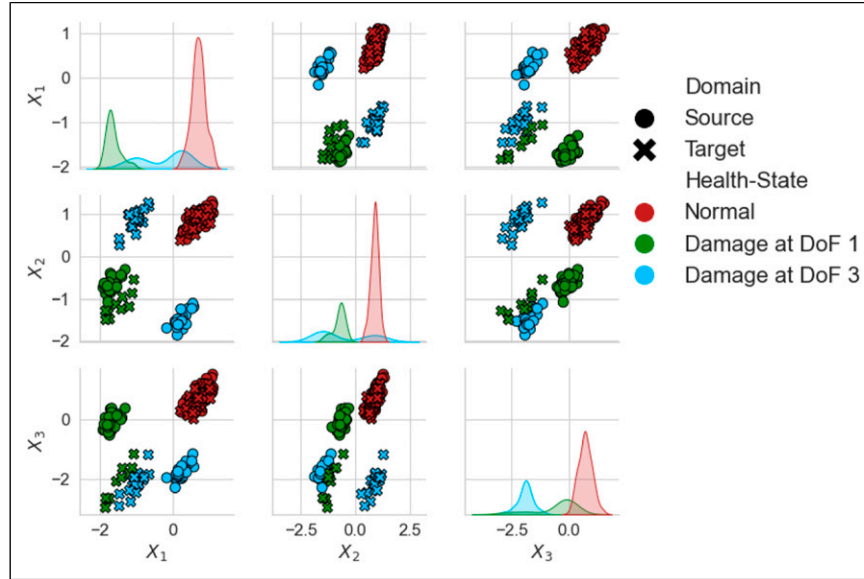


Figure 16. Normal-condition alignment features of the numerical three- to seven-storey population of structures. The source and target are depicted by (○) and (×), respectively. A random subset 20% the size of the datasets is plotted for visualisation.

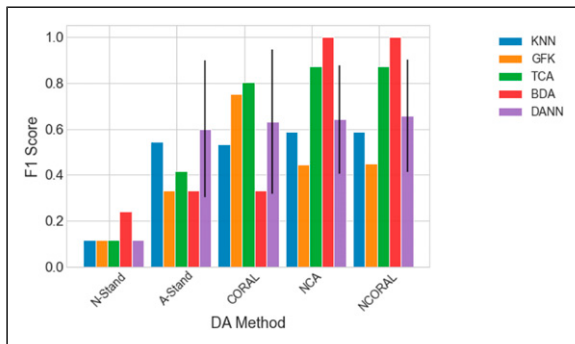


Figure 17. Classification performance of a KNN on the target domain after various SA methods, then DA on the numerical three- to seven-storey population. The result of the DANN is given as the mean of 100 repeats with one standard deviation shown by a black line. DANN: domain adversarial neural network; KNN: k-nearest neighbours classifier.

with the two structures having differences in material properties, geometries and connectivity. Thus, since the aim of DA here is to find a feature space where future health-state data could be shared, the features should share physical similarity, such that it is believed that the same physical phenomena in each structure would correspond to similar parts of the feature space. This objective justifies the approach of selecting only frequencies with strong modal correspondence, but the dissimilarity between the bridges means that only two frequencies were deemed to be sufficiently similar. In a population where structures are more similar, it is reasonable to expect that a larger proportion of modes would be similar, reducing the severity of this trade-off; for example, two nominally identical wind

turbines would be expected to have a larger proportion of similar modes.

Despite the masking effects, the aim of this case was to demonstrate that NCA and NCORAL can align the feature spaces of two structures in a partial DA scenario, as well as address changes to the structural response caused by repair by only using limited normal-condition data gathered in a short period after inspection. This objective has been achieved using a small quantity of inexpensive data, that is, normal-condition response and temperature data. The trade-off between transferable and damage sensitive features is a topic for further research.

Case study: SA as pre-processing

The previous case studies have demonstrated that exclusively aligning the lower-order statistics via SA can facilitate knowledge transfer. In fact, if the underlying data distributions are similar enough, SA can completely align the domains, as in the three-storey case study.

It is noted in contexts where the datasets are similar and SA is sufficient for adaptation, applying further adaptation may add an additional risk of negative transfer and lose the physical interpretability of the feature space. The question of whether similarity can be measured before transfer has been considered, with Refs. [4] and [30] proposing techniques for measuring physical and data similarity between members of a population.

In scenarios where members of the population are less similar, there may be disparity between both the lower- and the higher-order statistics. Thus, it is proposed that SA

should be used as pre-processing, aligning the lower-order statistics, such that the nonlinear mapping found by further DA is simplified.

Compared to SA, specifically NCA and NCORAL, prominent DA algorithms may have a number of additional requirements. Data available in each domain should be representative of their respective underlying distributions and abundant enough for nonparametric density estimation. Furthermore, prominent DA methods are prone to negative transfer in partial DA³⁶ and if the target only represents of small subset of the source classes, the risk of negative transfer will be higher.⁷¹ Therefore, further DA should be used when there is reason to assume that there are significant differences in the higher-order statistics and also, when available information is sufficient in each domain to mitigate the associated risk of negative transfer.

To demonstrate SA as a pre-processing method, another numerical case study is presented. The simulation procedure follows the previous numerical case study, with this case forming a heterogeneous population consisting of 3DoF source and 7DoF target structure, giving a more complex transfer problem (the material properties can be found in the [Appendix](#)). Damage was simulated for the first and third storeys in each structure. This case does not investigate partial DA, as prominent DA methods are prone to negative transfer in this scenario so identification of robust partial DA methods to use in conjunction with SA is left for future research. For the source, 400 samples were simulated for the normal-condition and 150 samples for each of the damage-states, $n_s = 700$. In the target, 200 samples were simulated for the normal-condition and 75 samples for each of the damage-states, $n_t = 350$. Class imbalance between the normal-condition and damage-states is introduced in this way to emulate practical scenarios, where normal-condition data will be more abundant. The TL problem is, therefore, to transfer the label information from the normal-condition, first damage location and third damage location in the source structure to the same location in the target structure.

N-standardisation, A-standardisation, CORAL and NCA were each used as a normalisation procedure. Following this step, each of the previously used DA methods were applied, including a KNN on the N-standardised space, with one neighbour, to provide a benchmark for performance gains given by further DA.

Results: Numerical three- to seven-storey population

The NCA features are visualised in [Figure 16](#). Significant differences in the domains can be seen, including the order of Classes 1 and 2 being flipped between domains in the second mode. Classification results can be found in [Figure 17](#). Without first performing SA, none of the nonlinear DA methods achieved adequate adaptation for knowledge transfer. Each of the SA methods alone

improved upon the unaligned KNN. Following NCA and NCORAL, excluding the GFK, all DA methods were able to improve upon SA alone. In this case, NCORAL gives the same result as NCA, since correlation of the normal-conditions is the same in both domains. The GFK did result in negative transfer after NCA and CORAL, but this method aligns PCA subspaces with dimension k , with the condition it must be under half the dimension of the original space d , that is, $k < d/2$. Thus, in this case, $k = 1$, which may have not been sufficient to encode enough discriminative information. A-standardisation and CORAL did perform as well as NCA or NCORAL with any method, with the exception of the GFK after CORAL, suggesting that class imbalance effected their performance. These results suggest that even though without SA, the nonlinear DA methods fail to achieve knowledge transfer, they can still provide additional performance gain if SA is applied first, highlighting that it is crucial to consider SA for transfer in SHM.

Conclusions

Population-based SHM aims to facilitate more in-depth health-state diagnosis in SHM by sharing information across a population of structures. When utilising information across different structures, whether that be a homogeneous or heterogeneous population, there will be differences between the training and testing distributions. Transfer learning, in the form of DA, offers a solution to this issue by finding a mapping that aligns the feature spaces. However, previously implemented DA methods have relied on methods that are data intensive, susceptible to negative transfer under class imbalance and in partial DA, and map data into a latent space that may make it challenging to interpret physical phenomena. SA is proposed as a solution to these challenges, with two methods proposed to address issues regarding class imbalance and partial DA. It is also proposed that SA can simplify the problem for nonlinear methods when used as a pre-processing method.

Three case studies were presented to evaluate the applicability of SA. The first, a heterogeneous numerical population, found that the differences between linear structural responses caused by material properties and the size of the structure could be removed by aligning the mean and standard deviations. Furthermore, the main approaches to DA that have been previously studied failed to adapt the domains. The dataset was then downsampled to investigate class imbalance, in a partial DA scenario, suggesting that previous SA approaches may cause negative transfer under class imbalance, whereas NCA and NCORAL are more robust.

The second case study presented a real heterogeneous population of two bridges – the Z24 and KW51 bridges. Normal-Correlation Alignment was applied, solving two partial DA problems relating to the pre- and post-repair

states in the KW51 bridge dataset, as well as adapting the KW51 and Z24 bridge datasets. Adaptation was achieved by using limited data directly after inspection in the KW51 bridge, which was shown to be adequate to identify behaviour of interest using an unsupervised GMM. In addition, the idea of sample selection using measurable EoVs to reduce class imbalance was demonstrated by selecting data corresponding to above and below freezing temperatures. Intuitively, this population may be considered as a highly challenging transfer scenario, as it consists of a concrete highway bridge and a steel railway bridge. Thus, the fact that NCORAL could align the domains with such a small dataset demonstrates the potential of aligning normal-condition data. A potential application of this alignment would be to define a more informed threshold for damage detection, using the damage in the Z24 bridge dataset. Such a threshold would result in no false positive in the KW51 bridge dataset despite KW51 bridge normal-condition data being extremely close to boundary for the identified damage cluster in the feature space after alignment.

In dissimilar populations, only considering the lower-order statistics may limit the potential performance gains that a nonlinear mapping could provide, so a final case study discusses using SA as a pre-processing method to simplify the DA problem. Even though the nonlinear DA methods utilised should have the ability to align the lower- and higher-order statistics, without first using SA, all of these methods failed to transfer knowledge. This result suggests that SA is a potentially crucial pre-processing method for DA in SHM. Normal-condition alignment and NCORAL were shown to be particularly robust to class imbalance.

Future work is required for the reliable application of the technologies discussed in this paper. One limitation of the DA identified here, is the trade-off between discriminative and transferable features, motivating the investigation of the required level of similarity between features for robust transfer. In addition, this paper suggests that for similar features, SA can achieve sufficient adaptation to provide more accurate; in-depth diagnostics and in situations where further DA is beneficial, SA is a crucial pre-processing step. Since there is a risk of negative transfer associated with applying DA,⁴¹ DA should not be naively applied, so future work is required to identify when it is beneficial and how to reduce the risk of negative transfer in an unsupervised manner. A large part of this work will be concerned with identifying a suitable unsupervised metric to measure similarity between features; for example, physical similarity may be used.⁴ Also, methods used in conjunction with SA for aligning higher-order statistics should also be robust to partial DA problems, so the development of partial DA algorithms that do not require large datasets will also be considered. Finally, partial DA has been demonstrated as a useful tool to alleviate the strict assumption that the label spaces are identical, as made by conventional DA.

Nevertheless, in some scenarios, the target domain will also contain data pertaining to a health-state that is not in the source, motivating the study of methods that further weaken this assumption; for example, open-DA.⁷

Acknowledgements

The authors would like to acknowledge the support of the UK Engineering and Physical Sciences Research Council via grants EP/R006768/1, EP/R003645/1 and EP/R004900/1. Lawrence Bull was supported by Wave one 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 providing access to the KW51 dataset.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by Engineering and Physical Sciences Research Council (EP/R003645/1, EP/R004900/1, EP/R006768/1, and EP/W006022/1).

ORCID iD

Jack Poole  <https://orcid.org/0000-0002-7642-9108>

Notes

1. It should be noted that these measurements will give an indication of what data are affected by certain EoVs, but they will not guarantee the data was generated by a given process.
2. The macro-F1 score for N-standardisation is higher than the previous case because there are only two classes in the target domain, so classifying all classes as one class results in a macro-F1 score of 0.5.
3. Even if there were more features available, the feature reduction performed in step one may make the problem a heterogeneous DA problem, where the feature spaces differ in dimension. Heterogeneous DA is considered more challenging and requires specialised technologies.⁴⁴
4. This assumption could also be enforced with Bayesian priors.
5. Note, the Z24 bridge damage labels were not used to learn this GMM, as the main aim is to show the domains are well aligned by SA. A supervised or semi-supervised model could be used to better define this boundary.

References

1. Farrar C and Worden K. *Structural Health Monitoring: A Machine Learning Perspective*. Chichester: Wiley, 2013.

2. Dervilis N, Antoniadou I, Barthorpe RJ, et al. Robust methods for outlier detection and regression for SHM applications. *Int J Sustain Mater Struct Syst* 2015; 2(1/2): 3.
3. Bull LA, Gardner P, Gosliga J, et al. Foundations of population-based SHM, Part I: Homogeneous populations and forms. *Mech Syst Signal Process* 2021; 148: 107141.
4. Gosliga J, Gardner P, Bull LA, et al. Foundations of Population-based SHM, Part II: heterogeneous populations – graphs, networks, and communities. *Mech Syst Signal Process* 2021; 148: 107144.
5. Gardner P, Bull L, Gosliga J, et al. Foundations of population-based SHM, Part III: Heterogeneous populations – mapping and transfer. *Mech Syst Signal Process* 2021; 149: 107142.
6. Bishop CM. *Pattern Recognition and Machine Learning*. New York: Springer, 2006.
7. Zhuang F, Qi Z, Duan K, et al. A comprehensive survey on transfer learning. *Proc IEEE* 2021; 109(1): 43–76.
8. Jialin Pan S, Tsang IW, Kwok JT, et al. Domain adaptation via transfer component analysis. *IEEE Trans Neural Netw* 2011; 22(2): 199–210.
9. Long M, Wang J, Ding G, et al. Transfer feature learning with joint distribution adaptation. In: Proceedings of the IEEE International Conference on Computer Vision, Sydney, NSW, Australia, 1–8 December 2013, pp. 2200–2207.
10. Long M, Wang J, Ding G, et al. Adaptation regularization: a general framework for transfer learning. *IEEE Trans Knowledge Data Eng* 2014; 26(5): 1076–1089. DOI: [10.1109/TKDE.2013.111](https://doi.org/10.1109/TKDE.2013.111).
11. Wang J, Chen Y, Hao S, et al. Balanced distribution adaptation for transfer learning. Proceedings - IEEE International Conference on Data Mining, ICDM 2017, New Orleans, LA, USA, 18 November 2017, pp. 1129–1134.
12. Long M, Cao Y, Wang J, et al. Learning transferable features with deep adaptation networks. 32nd International Conference on Machine Learning, ICML 2015, Lille, Nord, France, 1 June 2015, pp. 97–105.
13. Ganin Y and Lempitsky V. Unsupervised domain adaptation by backpropagation. International Conference on Machine Learning, Lille, Nord, France, 1 June 2015, pp. 1180–1189.
14. Yeh Ck, Matsuda N, Huang X, et al. Deep CORAL: Correlation Alignment for Deep Domain Adaptation. *European conference on computer vision*. Springer, Cham, 2016. Amsterdam, The Netherlands 2016; 9915: 738–752.
15. Tzeng E, Hoffman J, Saenko K, et al. Adversarial discriminative domain adaptation. In: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017. Honolulu, HI, USA, pp. 7167–7176.
16. Long M, Zhu H, Wang J, et al. Deep transfer learning with joint adaptation networks. In: 34th International Conference on Machine Learning, ICML 2017, PMLR, Sydney, Australia, 2017, pp. 3470–3479.
17. Long M, Cao Z, Wang J, et al. Conditional adversarial domain adaptation. *Adv Neural Inf Process Syst* 2018; 2018-Decem(NeurIPS): 1640–1650.
18. Zhao H, Zhang S, Wu G, et al. Adversarial multiple source domain adaptation. *Adv Neural Inf Process Syst* 2018; 2018-Decem(NeurIPS): 8559–8570.
19. Gong B, Shi Y, Sha F, et al. Geodesic flow kernel for unsupervised domain adaptation. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, 16–21 June 2012, pp. 2066–2073.
20. Chakraborty D, Kovvali N, Chakraborty B, et al. Structural damage detection with insufficient data using transfer learning techniques. *Sensors Smart Structures Tech Civil, Mech Aerospace Syst*. 2011; 7981: 798147. International Society for Optics and Photonics.
21. Xu S and Noh HY. Phymdan: Physics-informed knowledge transfer between buildings for seismic damage diagnosis through adversarial learning. *Mech Syst Signal Process* 2021; 151: 107374.
22. Wang X and Liu F. Triplet loss guided adversarial domain adaptation for bearing fault diagnosis. *Sensors* 2020; 20(1): 320.
23. Li X, Zhang W, Ding Q, et al. Multi-Layer domain adaptation method for rolling bearing fault diagnosis. *Signal Process* 2019; 157: 180–197.
24. Michau G and Fink O. Domain adaptation for one-class classification: monitoring the health of critical systems under limited information, 2019. arXiv preprint arXiv: 1907.09204.
25. Li Y, Song Y, Jia L, et al. Intelligent Fault Diagnosis by Fusing Domain Adversarial Training and Maximum Mean Discrepancy via Ensemble Learning. *IEEE Trans Ind Inform* 2020; 17(4): 2833–2841.
26. Jiao J, Zhao M, Lin J, et al. Residual joint adaptation adversarial network for intelligent transfer fault diagnosis. *Mech Syst Signal Process* 2020; 145: 106962.
27. Ding Y, Jia M and Cao Y. Remaining useful life estimation under multiple operating conditions via deep subdomain adaptation. *IEEE Trans Instrumentation Meas* 2021; 70: 1–11.
28. Gao Y and Mosalam KM. Deep transfer learning for image-based structural damage recognition. *Comput-Aided Civil Infrastruct Eng* 2018; 33(9): 748–768. DOI: [10.1111/mice.12363](https://doi.org/10.1111/mice.12363).
29. Dorafshan S, Thomas RJ and Maguire M. Comparison of deep convolutional neural networks and edge detectors for image-based crack detection in concrete. *Construction Building Mater* 2018; 186: 1031–1045.
30. Gardner P, Liu X and Worden K. On the application of domain adaptation in structural health monitoring. *Mech Syst Signal Process* 2020; 138: 106550.
31. Gardner P, Bull L, Gosliga J, et al. A population-based shm methodology for heterogeneous structures: Transferring damage localisation knowledge between different aircraft wings. *Mech Syst Signal Process* 2022; 172: 108918.

32. Gardner P, Bull L, Dervilis N, et al. Overcoming the problem of repair in structural health monitoring: Metric-informed transfer learning. *J Sound Vibration* 2021; 510: 116245.
33. Bull L, Gardner P, Dervilis N, et al. On the transfer of damage detectors between structures: An experimental case study. *J Sound Vibration* 2021; 501: 116072.
34. Song L and Dai B. Robust low rank kernel embeddings of multivariate distributions. In: *Advances in Neural Information Processing Systems*, Lake Tahoe, Nevada, USA, 2013, pp. 1–9.
35. Mao J, Wang H and Spencer BF Jr. Toward data anomaly detection for automated structural health monitoring: Exploiting generative adversarial nets and autoencoders. *Struct Health Monit* 2021; 20(4): 1609–1626.
36. Cao Z, Long M, Wang J, et al. Partial transfer learning with selective adversarial networks. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Salt Lake City, UT, USA, 2018, pp. 2724–2732.
37. Wang G, Sun T, Pan Y, et al. *Adaptive Neural Network Control for Constrained Robot Manipulators, Volume 10262 LNCS*, 2017.
38. Jiao J, Zhao M, Lin J, et al. Classifier inconsistency-based domain adaptation network for partial transfer intelligent diagnosis. *IEEE Trans Ind Inform* 2019; 16(9): 5965–5974.
39. Deng Y, Huang D, Du S, et al. A double-layer attention based adversarial network for partial transfer learning in machinery fault diagnosis. *Comput Industry* 2021; 127: 103399.
40. Zhang J, Li W, Ogunbona P, et al. Recent advances in transfer learning for cross-dataset visual recognition: A problem-oriented perspective. *ACM Comput Surv (Csur)* 2019; 52(1): 1–38.
41. Zhang W, Deng L and Wu D. *Overcoming Negative Transfer: A Survey*, 2020, pp. 1–15. *arXiv*.
42. Maeck J and De Roeck G. Description of Z24 benchmark. *Mech Syst Signal Process* 2003; 17(1): 127–131.
43. Maes K and Lombaert G. Monitoring railway bridge KW51 before, during, and after retrofitting. *J Bridge Eng* 2021; 26(3): 04721001.
44. Weiss K, Khoshgoftaar TM and Wang D. A survey of transfer learning. *J Big Data* 2016; 3(1): 1–40.
45. Sun Q, Chattopadhyay R, Panchanathan S and Ye J. A two-stage weighting framework for multi-source domain adaptation. *Advances in Neural Information Processing Systems* 2011; 24: 1–9.
46. Harel M and Mannor S. Learning from multiple outlooks. *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*, Bellevue, Washington, USA, 2011: 401–408.
47. He H and Wu D. Transfer learning for brain-computer interfaces: a euclidean space data alignment approach. *IEEE Trans Biomed Eng* 2020; 67(2): 399–410.
48. Li Y, Wang N, Shi J, et al. Adaptive Batch Normalization for practical domain adaptation. *Pattern Recognition* 2018; 80: 109–117.
49. Chang WG, You T, Seo S, et al. *Domain-specific Batch Normalization for Unsupervised Domain Adaptation*, 2019, pp. 7354–7362.
50. Seo S, Suh Y, Kim D, et al. Learning to optimize domain specific normalization for domain generalization. In: *Computer Vision—ECCV 2020: 16th European Conference, Proceedings, Part XXII 16*, Glasgow, UK, 23–28 August 2020. Springer, pp. 68–83.
51. Nam J, Pan SJ and Kim S. Transfer defect learning. *Proceedings - International Conference on Software Engineering* 2013, San Francisco, CA, USA, 18–26 May 2013, pp. 382–391.
52. Bull L, Worden K, Fuentes R, et al. Outlier ensembles: a robust method for damage detection and unsupervised feature extraction from high-dimensional data. *J Sound Vibration* 2019; 453: 126–150.
53. Dervilis N, Worden K and Cross E. On robust regression analysis as a means of exploring environmental and operational conditions for shm data. *J Sound Vibration* 2015; 347: 279–296.
54. Cross E, Koo K, Brownjohn J, et al. Long-term monitoring and data analysis of the tamar bridge. *Mech Syst Signal Process* 2013; 35(1–2): 16–34.
55. Christides S and Barr ADS. One-dimensional theory of cracked Bernoulli-Euler beams. *Int J Mech Sci* 1984; 26(11–12): 639–648.
56. Gardner P, Fuentes R, Dervilis N, et al. Machine learning at the interface of structural health monitoring and non-destructive evaluation. *Philosophical Trans R Soc A* 2020; 378(2182): 20190581.
57. van Baars Buisman S, Michau G, Rene A, et al. Domain adaptations for guided wave shm of composites: Towards fleet monitoring. *Proc 6th Eur Conf Prognostics Health Manage Soc* 2021; 6: 430–438. Prognostics and Health Management Society.
58. Gretton A, Sriperumbudur B, Sejdinovic D, et al. Optimal kernel choice for large-scale two-sample tests. *Adv Neural Inf Process Syst* 2012; 2: 1205–1213.
59. Goodfellow IJ, Pouget-Abadie J, Mirza M, et al. Generative adversarial nets. *Adv Neural Inf Process Syst* 2014; 3(-January): 2672–2680.
60. Poole J, Gardner P, Dervilis N, et al. On normalisation for domain adaptation in population-based structural health monitoring. To appear In *Proceedings of the 13th International Workshop on Structural Health Monitoring*, Stanford University, CA, 2021.
61. Murphy KP. *Machine Learning: A Probabilistic Perspective*. Cambridge, MA: MIT press, 2012.
62. Maeck J, Peeters B and De Roeck G. Damage identification on the Z24 bridge using vibration monitoring. *Smart Mater Struct* 2001; 10(3): 512.

63. Peeters B and De Roeck G. One-year monitoring of the Z24-bridge: environmental effects versus damage events. *Earthquake Eng Struct Dyn* 2001; 30(2): 149–171.
64. Roeck GD. The state-of-the-art of damage detection by vibration monitoring: the simces experience. *J Struct Control* 2003; 10(2): 127–134.
65. Teughels A and De Roeck G. Structural damage identification of the highway bridge Z24 by FE model updating. *J Sound Vibration* 2004; 278(3): 589–610.
66. Reynders E and De Roeck G. A local flexibility method for vibration-based damage localization and quantification. *J Sound Vibration* 2010; 329(12): 2367–2383.
67. Langone R, Reynders E, Mehrkanoon S, et al. Automated structural health monitoring based on adaptive kernel spectral clustering. *Mech Syst Signal Process* 2017; 90: 64–78.
68. Bull L, Rogers T, Wickramarachchi C, et al. Probabilistic active learning: An online framework for structural health monitoring. *Mech Syst Signal Process* 2019; 134: 106294.
69. Hughes A, Bull LA, Gardner P, et al. On risk-based active learning for structural health monitoring. *Mech Syst Signal Process* 2022; 167: 108569.
70. Maes K, Van Meerbeeck L, Reynders E, et al. Validation of vibration-based structural health monitoring on retrofitted railway bridge KW51. *Mech Syst Signal Process* 2022; 165: 108380.
71. Shui C, Chen Q, Wen J, et al. *Beyond -divergence: Domain Adaptation Theory with Jensen-Shannon Divergence*, 2020. *arXiv preprint arXiv:200715567*.

Appendix

Material properties

Table A1. Properties for the numerical 3-storey case study.

	Unit	Source	Target
Beam geometry $\{l_b, w_b, t_b\}$	mm	{300, 40, 8}	{160, 25, 6}
Mass geometry $\{l_m, w_m, t_m\}$	mm	{400, 400, 40}	{300, 250, 25}
Crack geometry $\{l_{cr}, l_{oc}\}$	mm	{20.0, 150}	{12.5, 80}
Elastic modulus E	GPa	$\mathcal{N}(210, 1 \times 10^{-9})$	$\mathcal{N}(71, 1 \times 10^{-10})$
Density ρ	kg/m ³	$\mathcal{N}(7800, 50)$	$\mathcal{N}(2700, 10)$
Damping coefficient c	Ns/m	$\mathcal{G}(8, 0.8)$	$\mathcal{G}(50, 0.8)$

Table A2. Properties for the three- to seven-storey numerical case study.

	Unit	Source	Target
Beam geometry $\{l_b, w_b, t_b\}$	mm	{300, 40, 8}	{300, 40, 8}
Mass geometry $\{l_m, w_m, t_m\}$	mm	{400, 400, 40}	{400, 400, 40}
Crack geometry $\{l_{cr}, l_{oc}\}$	mm	{20.0, 150}	{20.0, 150}
Elastic modulus E	GPa	$\mathcal{N}(210, 1 \times 10^{-9})$	$\mathcal{N}(210, 1 \times 10^{-9})$
Density ρ	kg/m ³	$\mathcal{N}(7800, 50)$	$\mathcal{N}(7800, 50)$
Damping coefficient c	Ns/m	$\mathcal{G}(8, 0.8)$	$\mathcal{G}(8, 0.8)$
Storeys		3	7